arXiv:1606.02703v2 [math.PR] 27 Jun 2016

Mixing time for exclusion and interchange processes on hypergraphs

Stephen B. Connor* and Richard J. Pymar†

28th June 2016

Abstract

We introduce a natural extension of the exclusion process to hypergraphs and prove an upper bound for its mixing time. In particular we show the existence of a constant C such that for any connected hypergraph G within some class (which includes regular hypergraphs), the ε-mixing time of the exclusion process on G with any feasible number of particles can be upper-bounded by $CT_{\mathrm{EX}(2,G)}\log(|V|/\varepsilon)$, where $|V|$ is the number of vertices in G and $T_{\mathrm{EX}(2,G)}$ is the 1/4-mixing time of the corresponding exclusion process with two particles. We also obtain a similar result for the interchange process provided the number of particles is at most $|V|/2$. The proofs involve an adaptation of the *chameleon process*, a technical tool invented by Morris (2006) and developed by Oliveira (2013) for studying the exclusion process on a graph.

Keywords and phrases:
MIXING TIME, EXCLUSION, INTERCHANGE, RANDOM WALK, HYPERGRAPH, COUPLING

2000 Mathematics Subject Classification:
Primary: 60J27, 60K35
Secondary: 82C22

1 Introduction

Let $G = (V, E)$ be a finite connected graph with vertex set V and edge set E. Fix $k \in \{1, \ldots, |V|\}$ and consider k indistinguishable particles moving on V using the following rules:

1. each vertex is occupied by at most one particle,

*Work supported in part by EPSRC Research Grant EP/J009180.
†Work supported in part by Leverhulme Research Grant RPG-2012-608 held by Nadia Sidorova and EPSRC Grant EP/M027694/1 held by Codina Cotar.

2. each edge $e \in E$ rings at the times of a Poisson process of rate 1, independently of all other edges,

3. when an edge $e = \{u, v\}$ rings, the occupancy states of vertices u and v are switched.

For each $v \in V$ and $t \geq 0$, let $\eta_t(v) = 1$ if v is occupied at time t, and $\eta_t(v) = 0$ if v is vacant at time t. The process $(\eta_t)_{t\geq 0}$ is called the *k-particle exclusion process on* G: see Figure 1. In this paper we are interested in a natural extension of the exclusion process to *hypergraphs*.

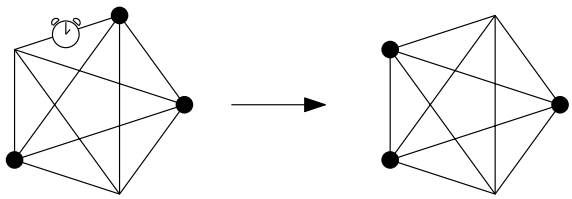

Figure 1: Example transition of 3-particle exclusion process on K_5. When the edge indicated rings, the single particle currently on that edge moves to the vertex at the other end of the edge.

Let $G = (V, E)$ be a finite connected hypergraph, where $E \subseteq \mathcal{P}(V)$, the power set of V. For each $e \in E$, denote by $\mathcal{S}_e$ the symmetric group of e, and let $f_e : \mathcal{S}_e \to [0, 1]$ be a probability measure on $\mathcal{S}_e$. We write f to denote $\{f_e : e \in E\}$, the set of these measures. Consider k indistinguishable particles moving on V using rules 1. and 2. above and in addition:

3′. when an edge e rings, a permutation $\sigma \in \mathcal{S}_e$ is chosen according to f_e and every particle on a vertex in e moves simultaneously according to σ, i.e. a particle at vertex v moves to vertex $\sigma(v)$ (note that as σ is a permutation, rule 1. is preserved).

Setting $\eta_t(v) = 1$ if v is occupied at time t and 0 otherwise, we obtain a process $(\eta_t)_{t\geq 0}$ referred to as the *k-particle exclusion process on* (f, G), or simply $\mathrm{EX}(k, f, G)$: see Figure 2. (Note that if each edge $e \in E$ contains exactly two vertices, and f_e puts all of its mass on the transposition belonging to $\mathcal{S}_e$, then $\mathrm{EX}(k, f, G)$ is just the k-particle exclusion process on the graph G, as above.)

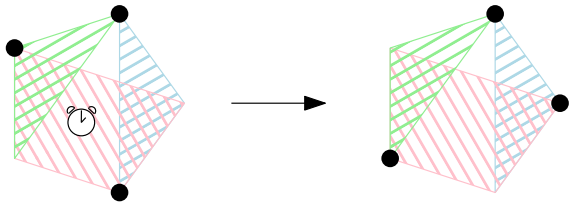

Figure 2: Example transition of 3-particle exclusion process on a hypergraph with 5 vertices and 3 edges (indicated by the different shaded regions). When the edge containing four vertices rings, the two particles currently belonging to that edge are permuted.

Our main aim in this paper is to study the total-variation mixing time of $\mathrm{EX}(k, f, G)$, and to establish an upper bound in terms of the mixing time of $\mathrm{EX}(2, f, G)$. Recall that for a continuous-time Markov process X on a finite set Ω with transition probabilities $\{q_t(x, y)\}$ and equilibrium distribution π, the total variation ε-mixing time is defined as

$$T_X(\varepsilon) := \inf \left\{ t \geq 0 : \max_{x\in\Omega} \|q_t(x, \cdot) - \pi\|_{\mathrm{TV}} \leq \varepsilon \right\}, \tag{1.1}$$

where $\|\cdot\|_{\mathrm{TV}}$ is the total-variation norm. The exclusion process (on either a connected graph or a connected hypergraph) is irreducible with symmetric transition rates, and therefore has a uniform stationary distribution.

In addition to studying the mixing time of $\mathrm{EX}(k, f, G)$ we will also study the mixing time of the equivalent process where the k particles are distinguishable. Suppose the particles are labelled $1, \ldots, k$ and set $\hat{\eta}_t(v)$ to be the label of the particle at vertex v at time t. If there is no particle at v at time t, set $\hat{\eta}_t(v) = 0$. The process $(\hat{\eta}_t)_{t\geq 0}$ is the *k-particle interchange process on* (f, G), or simply $\mathrm{IP}(k, f, G)$.

We make the following assumptions about the hypergraph G and set of measures f (notation appearing below is defined in Section 2.1).

Assumption 1.1 (Assumptions on G)**.**

1. G is a connected hypergraph, with $\min_{e\in E} |e| \geq 7$ and $|V| \geq 36$;
2. the following bound holds:

$$\sum_{e\in E} \sum_{a,b\in\binom{e}{2}} \frac{2}{(|e|-1)(\deg(a)+\deg(b))} \geq \frac{|V|}{10}\,, \tag{1.2}$$

   where, for $v \in V$, $\deg(v)$ is the number of edges containing v.

Assumption 1.2 (Assumptions on f)**.**

1. For every e, f_e is constant on the conjugacy classes of $\mathcal{S}_e$ (i.e. in group-theoretic terms, f_e is a *class function*). That is, if σ_1 and σ_2 are elements from $\mathcal{S}_e$ with the same cycle structure, then $f_e(\sigma_1) = f_e(\sigma_2)$.
2. For every e and each $v \in e$, $\sum_{\sigma\in\mathcal{S}_e} f_e(\sigma)\mathbf{1}_{\{|\sigma(v)|\leq 3\}} \leq 1/4$, where $|\sigma(v)|$ denotes the size of the cycle containing v in σ. In other words, the probability (under f_e) of a vertex $v \in e$ being in a cycle of σ of size less than four is at most $1/4$.

Our main theorem is the following:

Theorem 1.3. *There exists a universal constant $C > 0$ such that for every hypergraph G satisfying Assumption 1.1, every f satisfying Assumption 1.2 and every $k \in \{1, \ldots, |V| - 1\}$ and $\varepsilon > 0$,*

$$T_{\mathrm{EX}(k,f,G)}(\varepsilon) \leq C \log(|V|/\varepsilon) T_{\mathrm{EX}(2,f,G)}(1/4).$$

Our second main result relates the mixing time of the k-particle interchange process to the mixing time of a 2-particle exclusion process:

Theorem 1.4. *There exists a constant $C > 0$ such that for every hypergraph G satisfying Assumption 1.1, every f satisfying Assumption 1.2 and every $k \in \{1, \ldots, |V|/2\}$ and $\varepsilon > 0$,*

$$T_{\mathrm{IP}(k,f,G)}(\varepsilon) \leq C \log(|V|/\varepsilon) T_{\mathrm{EX}(2,f,G)}(1/4).$$

Remark 1.5. The fraction 1/10 appearing in Part 2 of Assumption 1.1 is arbitrary in the sense that it can be set to any other constant, the only effect being a change in the universal constant C. An alternative formulation for this condition is the following: let $d^* = \sum_{v \in V} \deg(v) = \sum_{e \in E} |e|$, and $\bar{d} = d^*/|V|$. Then (1.2) is equivalent to requiring

$$\sum_{e \in E} \frac{|e|}{d^*} \sum_{a,b \in \binom{e}{2}} \frac{\binom{|e|}{2}^{-1}}{\deg(a) + \deg(b)} \geq 1/10\bar{d}\,. \tag{1.3}$$

That is, suppose that we randomly choose an edge $e \in E$ with probability proportional to $|e|$. Given e, uniformly choose a pair of vertices $a, b \in e$, and compute $1/(\deg(a) + \deg(b))$. The result of doing this should be bounded below by the reciprocal of ten times the mean vertex degree, $\bar{d}$.

A simple application of Jensen's inequality shows that the left hand side of (1.3) is bounded below by $d^*/(2\sum_{v \in V} \deg(v)^2)$. It follows that a sufficient condition for Part 2 of Assumption 1.1 to hold is that

$$\mathbb{E}[\deg(v)]^2 \geq \frac{1}{5}\,\mathbb{E}[\deg(v)^2]\,,$$

where v is a uniformly chosen vertex in the hypergraph.

Inequality (1.2) is clearly satisfied by a regular hypergraph (in which all vertices have equal degree) and, for example, by a d-star hypergraph (in which there are $(|V|-1)/(d-1)$ edges, all of size d, with one central vertex belonging to every edge and no other vertex belonging to more than one edge: see Figure 3). On the other hand, consider modifying the d-star by having a set of $(d-1)$ 'central' vertices that belong to every edge. (That is, every edge consists of these central vertices plus one other non-central vertex.) This no longer satisfies (1.2).

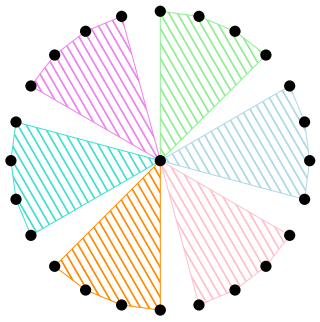

Figure 3: 5-star on 25 vertices

Remark 1.6. Our analysis can be extended to *weighted* hypergraphs, in which each edge $e \in E$ rings at some rate $w_e > 0$. Indeed, the only real change required is a modification of (1.2): the appropriate inequality becomes

$$\sum_{e \in E} \sum_{a,b \in \binom{e}{2}} \frac{2w_e}{(|e|-1)(\deg_w(a) + \deg_w(b))} \geq \frac{|V|}{10}\,,$$

where, for $v \in V$, $\deg_w(v) = \sum_{e : v \in e} w_e$.

In addition, we believe that Part 1 of Assumption 1.1 and Part 2 of Assumption 1.2 can be relaxed without changing the validity of our main results; however, it would appear that significantly more intricate arguments are needed to deal with such settings. In order to keep the presentation as simple as possible we have chosen to omit these, and to work under the assumption of equal weights on all edges.

1.1 Motivation and connections with the literature

Our results contribute to the general question of when properties of a multi-particle system can be deduced from properties of a system with only a few particles. Arguably the most significant recent result in this area has come from Caputo, Liggett, and Richthammer (2010) who showed that the spectral gap of the interchange process on a graph is equal to the spectral gap of a random walker on the same graph, proving a conjecture of Aldous that had been open for 20 years. Providing results in this area is particularly important in applications due to the large reduction in the size of the state space which often makes it much easier to compute or estimate statistics.

While interacting particle system models (e.g. exclusion process, interchange process, voter model, contact process, zero range process) on graphs has received considerable attention, there has so far been little study of such processes on hypergraphs. Studying these processes on hypergraphs is very natural though, as hypergraphs allow simultaneous interactions of multiple particles, rather than simply pair-wise interactions. One model for which its analogue on hypergraphs has been recently studied is the voter model (Chung and Tsiatas (2014); Istrate, Bonchis, and Marin (2014)), for which various properties are considered, including the mixing time.

Any interchange process (with $k = |V|$) on a graph can be viewed as a card shuffle by transpositions, with notable examples being the top-to-random transposition shuffle (star graph), random-to-random transposition shuffle (complete graph) and nearest-neighbour shuffle (the cycle). This connection is well-known, see e.g. Section 4A of Diaconis and Saloff-Coste (1993). The literature concerning mixing times of card shuffles by transpositions is extensive. In the following discussion we suppose that time is scaled so that shuffles are performed at unit rate (which amounts to setting the rate that each edge rings to $1/|E|$).

The case of the top-to-random transposition shuffle has been dealt with by Flatto, Odlyzko, and Wales (1985) and Diaconis (1988) (among others) – this shuffle has mixing time cutoff at time $|V| \log |V|$. The mixing time of the random-to-random transposition shuffle was first obtained by Diaconis and Shahshahani (1981), and is $\frac{|V|}{2} \log |V|$ (with cutoff). More recently, Lacoin (2016) has studied the nearest-neighbour transposition shuffle and showed mixing (with cutoff) occurs at time $\pi^{-2}|V|^3 \log |V|$. Jonasson (2012) has shown that the slowest a transposition shuffle can mix is of order $|V|^4 \log |V|$, with the lollipop graph an example which achieves this (Erikshed (2011)). Achieving these results typically involves finding an argument tailored specifically to the model in question. If we care less about the specific constant multiple (at which mixing occurs) and instead focus on the order, a result of Oliveira (2013) can prove particularly useful as a general way of bounding mixing times of interchange and exclusion processes (note however that it does not include the case $k = |V|$ for interchange):

Theorem 1.7 (Oliveira (2013)). *There exists a constant $C > 0$ such that for every connected weighted graph G and every $k \in \{1, \dots, |V| - 1\}$ and $\varepsilon \in (0, 1/2)$,*

$$T_{\mathrm{EX}(k,G)}(\varepsilon) \leq C \log(|V|/\varepsilon) T_{\mathrm{RW}(G)}(1/4)\,,$$

where $T_{\mathrm{RW}(G)}(1/4)$ is the mixing time of the random walk on G. Furthermore, for every $k \in \{1, \dots, |V|/2\}$

$$T_{\mathrm{IP}(k,G)}(\varepsilon) \leq C \log(|V|/\varepsilon) T_{\mathrm{RW}(G)}(1/4)\,.$$

Of course, transposition shuffles are just one class of shuffle, and there is significant interest in mixing times of more general shuffles in which multiple cards are moved simultaneously. A large class of time-homogeneous shuffles can be represented as interchange processes on hypergraphs. One of the simplest hypergraphs to consider is the complete uniform hypergraph (in which every edge contains the same number of vertices, and every possible such edge is in the hypergraph). If each edge has size d and if, when an edge rings, we move particles by application of a uniformly chosen cycle, this corresponds to shuffling cards by d-cycles. The mixing time (with cutoff) is known to be $\frac{|V|}{d} \log |V|$. This was first shown for the case of fixed d by Berestycki, Schramm, and Zeitouni (2011), and later extended to any $d = o(|V|)$, see Hough (2015). A related model is the following: instead of choosing a cycle uniformly, we fix a conjugacy class Γ (so Γ is a set of permutations with the same cycle structure) with $|\Gamma| = d$ (where $|\Gamma|$ is the number of non-fixed points in any permutation from Γ), and every time an edge rings, we choose a permutation on the vertices in that edge uniformly from Γ. Sharp bounds (up to a multiplicative constant) on this mixing time of this model with $d = o(|V|)$ were first obtained by Roichman (1996) with cutoff at time $\frac{|V|}{d} \log |V|$ proven by Berestycki and Sengul (2014). Cutoff at the same time for the model with $d \geq |V|/2$ was shown by Lulov and Pak (2002).

Our main results (Theorems 1.3 and 1.4) extend Theorem 1.7 to a class of hypergraphs; example applications are sketched in Section 1.2. Furthermore, our results hold for a large class of measures acting on the symmetric group $\mathcal{S}_{|V|}$ which goes beyond the standard framework studied by previous authors, in which a conjugacy class is fixed and then sampled from uniformly. Indeed, our measures f_e can vary dramatically between edges $e \in E$, and furthermore we do not require f_e to be supported on a fixed conjugacy class.

On the other hand, the conclusions of Theorems 1.3 and 1.4 are a little weaker than we would like: ideally the upper bounds would involve the mixing time of a simple random walk on (f, G) (as in Theorem 1.7), and we conjecture this stronger result to be true. (Note that this conjecture depends crucially upon f_e being constant on conjugacy classes. The running example in Diaconis and Saloff-Coste (1993), for which this condition is not satisfied, provides an example where the mixing time of two particles is at least $|V|$ times that of a single particle.) As will be explained below, the main obstacle to moving from the mixing time of $\mathrm{EX}(2, f, G)$ to that of $\mathrm{RW}(f, G)$ is the failure of negative correlation between the two processes when G is a hypergraph. We also conjecture that the regularity condition in part 2 of Assumption 1.1 can be relaxed.

1.2 Applications

In order to apply Theorems 1.3 and 1.4, to specific examples we just need to bound the mixing time of two particles moving under exclusion on the given hypergraph. We present a rough

outline demonstrating how Theorem 1.3 can be used for three different graphs. In each case we consider the k-particle exclusion process with $k = |V|/2$ and suppose that for each $e \in E$, f_e is zero on cycles of size not equal to $|e|$ and is constant on cycles of size $|e|$ (and thus satisfies Assumption 1.2 provided $|e| \geq 4$).

1. G is the d-uniform complete hypergraph, with all edges of size $d = o(|V|)$: a simple coupling of two copies of $\mathrm{EX}(2, f, G)$ gives an upper-bound on the total variation distance (using (2.4)) at time t of $2e^{-c_1 \lambda t}$, for some positive c_1 (not depending on $|V|$) where $\lambda = \binom{|V|-1}{d-1}$. This is less than $1/4$ provided $t > \frac{c_2}{\lambda}$, for some constant c_2. Theorem 1.3 gives us the bound
$$T_{\mathrm{EX}(k,f,G)}(\varepsilon) \leq \frac{c_3}{\lambda} \log(|V|/\varepsilon),$$
for some constant c_3. A lower bound of the same order can be obtained by considering the number of initially occupied vertices that are still occupied at time t under the exclusion process. Denoting by S the subset of V corresponding to the initially occupied vertices (note that $|S| = |V|/2$), one can consider the event A that there exists a subset of S of size at least $|V|/4 + \sqrt{|V|}$ containing only occupied vertices. Under the uniform measure, the probability of A tends to 0 as $|V| \to \infty$, whereas the probability of A in an exclusion process at time t is at least ε for $t < \frac{c_4}{\lambda} \log(|V|/\varepsilon)$, for a constant c_4. This shows that the upper bound obtained via Theorem 1.3 gives the correct order.

2. G is the d-star hypergraph described in Remark 1.5 (with $d = o(|V|)$): again, a coupling argument can be used for $\mathrm{EX}(2, f, G)$ to obtain an upper bound for the total variation distance from uniform at time t of $2e^{-c_1 t/d}$ for some positive constant c_1; a lower bound can be found by again considering the number of initially occupied vertices that are still occupied at time t. We find that Theorem 1.3 again gives the correct order of the mixing time, which is $d \log(|V|)$.

3. G is the hypergraph obtained from the d-star by replacing the central vertex with $|V|/d$ vertices, each of degree 2: see Figure 4. We can couple two copies of $\mathrm{EX}(2, f, G)$ (by waiting for pairs of particles to occupy the central edge, at which point we can couple them) to obtain an upper bound for the total variation distance from uniform at time t of $2e^{-c_1 t/d}$. A lower bound can be found (using the same argument as in the previous two examples) to give the order of the mixing time as $d \log(|V|)$.

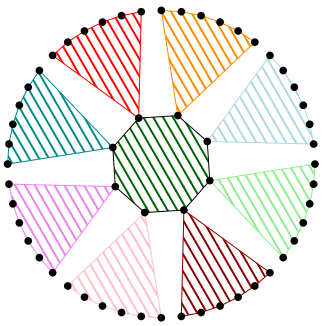

Figure 4: The central vertex in the 7-star (with $|V| = 56$) is replaced with 8 vertices of degree 2.

1.3 Heuristics and structure of the proof

In the rest of the paper we focus on proving Theorem 1.4; Theorem 1.3 actually follows as a corollary. To see this, one can simply 'forget' the labels of the particles in the interchange process to obtain the exclusion process and from this it is clear that once the interchange process is mixed, so is the exclusion process (this can be made rigorous using the contraction principle of Aldous and Fill (2002)).

The proof of Theorem 1.4 is split into two parts and follows in structure the work of Oliveira (2013). Indeed, in that paper the author proves his main result (bounding the mixing time of the k-particle interchange process by the mixing time of a random walker) by first relating the mixing time of a k-particle interchange process to that of a 2-particle interchange process. Roughly speaking, this is possible due to the fact that any time an edge of the graph under consideration rings, at most two particles move under interchange, and so it is pairwise interactions that determine the mixing rate. This contrasts with the interchange process on hypergraphs considered here, in which *many* particles can move at the same time. Nevertheless, a suitable adaptation of the techniques appearing in Oliveira (2013) yield a relationship between the mixing time of the k-particle interchange process and that of the 4-particle exclusion process:

Lemma 1.8. *There exists a constant $C > 0$ such that for every hypergraph G satisfying Assumption 1.1, every f satisfying Assumption 1.2 and every $k \in \{1, \dots, |V|/2\}$ and $\varepsilon > 0$,*

$$T_{\mathrm{IP}(k,f,G)}(\varepsilon) \leq C \log(|V|/\varepsilon) T_{\mathrm{IP}(4,f,G)}(1/4).$$

The second part of the proof requires showing that $T_{\mathrm{IP}(4,f,G)}$ and $T_{\mathrm{EX}(2,f,G)}$ are of the same order. We carry out the necessary calculations for this in Section 3.

Lemma 1.9. *For any hypergraph G satisfying Assumption 1.1 and f satisfying Assumption 1.2,*

$$T_{\mathrm{IP}(4,f,G)}(1/4) \leq 10^{14} T_{\mathrm{EX}(2,f,G)}(1/4).$$

We now demonstrate that our two main theorems follow simply from these two lemmas.

Proof of Theorem 1.4. This follows trivially by combining Lemma 1.9 and 1.8. □

Proof of Theorem 1.3. The contraction principle (see Aldous and Fill (2002)) gives

$$T_{\mathrm{EX}(k,f,G)}(\varepsilon) \leq T_{\mathrm{IP}(k,f,G)}(\varepsilon),$$

and so provided $k \leq |V|/2$, we have the result. However, note that $\mathrm{EX}(k, f, G)$ and $\mathrm{EX}(|V| - k, f, G)$ are the same process except roles of unoccupied and occupied vertices have been switched. It follows that $T_{\mathrm{EX}(k,f,G)}(\varepsilon) = T_{\mathrm{EX}(|V|-k,f,G)}(\varepsilon)$, which completes the proof. □

We finish this section with a brief overview of the rest of the paper. In Section 2 we define formally the processes considered in this paper and present some preliminary results; we also demonstrate that the negative correlation property, which is fundamental to the result in Oliveira (2013), fails to hold for the hypergraph setting. In Section 3 we prove Lemma 1.9 by

first characterizing every graph as one of two types depending on how long it takes any two of four particles moving under exclusion to 'feel' the presence of each other. Heuristically, if this time is likely to be of order $T_{\mathrm{EX}(2,f,G)}$ then, since the four particles essentially behave independently up to this time, they should be close to mixed. For graphs in which this is not the case, i.e. any two of four particles meet quickly, we adapt a coupling argument of Aldous and Fill (2002) to prove the lemma.

The most challenging (and technical) part of the proof is Lemma 1.8. The key component is an understanding of the movement of the kth particle in an interchange process when we condition on the trajectory of the other $k-1$ particles. Loosely speaking, we do this in the following way: firstly, we shall colour the $k-1$ other particles all *black*, the kth particle *red*, and we shall place a *white* particle onto each empty vertex. We shall then consider the amount of *redness* or (as we shall later refer to it) *ink* at each vertex, where initially only the vertex occupied by the red particle has any ink. In order to capture how mixed the kth particle is, we shall construct a process (later referred to as the *chameleon process*, and based on a simpler process of the same name appearing in Morris (2006) and Oliveira (2013)) in which ink is able to spread among the non-black particles in such a way that the faster it is able to spread, the faster the kth particle will mix. In Section 4 we introduce our version of the chameleon process and highlight some of its properties. Finally, in Sections 5 and 6 we use the chameleon process to prove Lemma 1.8.

2 Preliminaries

2.1 Random walks, exclusion and interchange processes

We formally define the main processes studied in this paper, $\mathrm{RW}(f,G)$, $\mathrm{RW}(k,f,G)$, $\mathrm{EX}(k,f,G)$ and $\mathrm{IP}(k,f,G)$, by explicitly stating their generators. In the next section we shall present a *graphical construction* of these processes, similar to that of Liggett (1999) for the standard interchange and exclusion processes. This graphical construction will allow us to simultaneously define the processes on the same probability space, and thus directly compare them.

Recall $\mathcal{S}_e$ as the group of permutations of elements in e. Our processes of interest evolve by the action of permutations from these groups. However, it will often be convenient to consider permutations as acting on V and we can easily do this by extending a permutation $\sigma_e \in \mathcal{S}_e$ to a permutation in $\mathcal{S}_V$ by setting $\sigma_e(v) = v$ for all $v \notin e$. We can also consider such permutations as acting on a subset of V or on vectors with elements being distinct members of V. To do this we can define, for a set $A \subseteq V$, $\sigma_e(A) := \{\sigma_e(a) : a \in A\}$, and for a vector $\mathbf{x}$ of k distinct elements of V we define $\sigma_e(\mathbf{x}) := (\sigma_e(\mathbf{x}(i)))_{i=1}^k$.

Set notation: For $k \in \mathbb{N}$ we define

$$\binom{V}{k} := \{A \subseteq V : |A| = k\},$$

and for a set $A \subseteq V$ we write

$$(A)_k := \{\mathbf{a} = (\mathbf{a}(1), \ldots, \mathbf{a}(k)) \in A^k : \mathbf{a}(i) \neq \mathbf{a}(j)\, \forall i \neq j\}.$$

Generators: We now explicitly state the generators of the processes. For a hypergraph G and a suitable set of functions f, the simple random walk on G, $\mathrm{RW}(f, G)$, is the continuous-time Markov chain with state space V and generator

$$U^{\mathrm{RW}}h(u) = \sum_{v\in V}\sum_{e\in E}\sum_{\substack{\sigma\in\mathcal{S}_e:\\ \sigma(u)=v}} f_e(\sigma)(h(v)-h(u))$$

for all $u \in V$ and $h : V \to \mathbb{R}$.

We denote by $\mathrm{RW}(k, f, G)$ the product of k independent random walkers on G. This process is the continuous-time Markov chain with state space V^k and generator

$$U^{\mathrm{RW}(k)}h(\mathbf{u}) = \sum_{v\in V}\sum_{e\in E}\sum_{i=1}^{k}\sum_{\substack{\sigma\in\mathcal{S}_e:\\ \sigma(\mathbf{u}(i))=v}} f_e(\sigma)(h(\mathbf{u}^i_v)-h(\mathbf{u})),$$

for all $\mathbf{u} \in V^k$ and $h : V^k \to \mathbb{R}$, where

$$\mathbf{u}^i_v(j) = \begin{cases}\mathbf{u}(j) & j\neq i,\\ v & j=i.\end{cases}$$

The k-particle exclusion process $\mathrm{EX}(k, f, G)$, is the continuous-time Markov chain with state space $\binom{V}{k}$ and generator

$$U^{\mathrm{EX}}h(A) = \sum_{B\in\binom{V}{k}}\sum_{e\in E}\sum_{\substack{\sigma\in\mathcal{S}_e:\\ \sigma(A)=B}} f_e(\sigma)(h(B)-h(A)),$$

for all $A \in \binom{V}{k}$ and $h : \binom{V}{k} \to \mathbb{R}$.

The k-particle interchange process $\mathrm{IP}(k, f, G)$, is the continuous-time Markov chain with state space $(V)_k$ and generator

$$U^{\mathrm{IP}}h(\mathbf{x}) = \sum_{\mathbf{y}\in(V)_k}\sum_{e\in E}\sum_{\substack{\sigma\in\mathcal{S}_e:\\ \sigma(\mathbf{x})=\mathbf{y}}} f_e(\sigma)(h(\mathbf{y})-h(\mathbf{x})),$$

for all $\mathbf{x} \in (V)_k$ and $h : (V)_k \to \mathbb{R}$.

2.2 Graphical construction

We first construct an independent sequence of E-valued random variables $\{e_n\}_{n\in\mathbb{N}}$ such that each e_n is identically distributed with $\mathbb{P}[e_n = e] = 1/|E|$. Given the sequence $\{e_n\}_{n\in\mathbb{N}}$, let $\{\sigma_n\}_{n\in\mathbb{N}}$ be a sequence of permutations with $\sigma_n \in \mathcal{S}_{e_n}$ independently chosen and satisfying for each $n \in \mathbb{N}$, $\mathbb{P}[\sigma_n = \sigma] = f_{e_n}(\sigma)$. Now that we have the sequence of edges that ring and the permutations to apply, it remains to determine the update times of the processes.

Let Λ be a Poisson process of rate $|E|$ and for $0 < s < t$ denote by $\Lambda[s, t]$ the number of points of Λ in $[s, t]$. For every $0 < s < t$, we define a random permutation $I_{[s,t]} : V \to V$ associated

with the time interval $[s,t]$ to be the composition of the permutations performed during this time; that is,

$$I_{[s,t]} = \sigma_{e_{\Lambda[0,t]}} \circ \sigma_{e_{\Lambda[0,t]-1}} \circ \cdots \circ \sigma_{e_{\Lambda[0,s)+1}}.$$

We set $I_t := I_{[0,t]}$ for each $t > 0$, and $I_{(t,t]}$ to be the identity. We can lift the functions $I_{[s,t]}$ to functions on $\binom{V}{k}$ and $(V)_k$ in the following way: for $A \in \binom{V}{k}$,

$$I_{[s,t]}(A) = \{I_{[s,t]}(a) : a \in A\},$$

and for $\mathbf{x} \in (V)_k$,

$$I_{[s,t]}(\mathbf{x}) = (I_{[s,t]}(\mathbf{x}(1)), \ldots, I_{[s,t]}(\mathbf{x}(k))).$$

The following proposition is fundamental and the proof follows by inspection.

Proposition 2.1. *Fix $s > 0$. Then:*

1. *For each $u \in V$, the process $\{I_{[s,s+t]}(u)\}_{t\geq 0}$ is a realisation of $\mathrm{RW}(f,G)$ initialised at u at time s. We shall often write this process simply as $(u_t^{\mathrm{RW}})_{t\geq s}$.*
2. *For each $A \in \binom{V}{k}$, the process $\{I_{[s,s+t]}(A)\}_{t\geq 0}$ is a realisation of $\mathrm{EX}(k,f,G)$ initialised at A at time s. We shall often write this process simply as $(A_t^{\mathrm{EX}})_{t\geq s}$.*
3. *For each $\mathbf{x} \in (V)_k$, the process $\{I_{[s,s+t]}(\mathbf{x})\}_{t\geq 0}$ is a realisation of $\mathrm{IP}(k,f,G)$ initialised at $\mathbf{x}$ at time s. We shall often write this process simply as $(\mathbf{x}_t^{\mathrm{IP}})_{t\geq s}$.*

2.3 Total variation and mixing times

There are several equivalent definitions of total variation that we shall make use of in this paper. Suppose μ and ν are two probability measures on the same finite set Ω. Then the *total variation distance* between these measures is defined as

$$\|\mu - \nu\|_{\mathrm{TV}} := \max_{A\subset\Omega}(\mu(A) - \nu(A)) \tag{2.1}$$

$$= \sup_{f:\Omega\to[0,1]} \int f d\mu - \int f d\nu \tag{2.2}$$

$$= \frac{1}{2}\sum_{x\in\Omega} |\mu(x) - \nu(x)|. \tag{2.3}$$

We shall also make extensive use of the following equivalent definition, which relates the total variation distance to couplings of μ and ν:

$$\|\mu - \nu\|_{\mathrm{TV}} = \inf_{(X,Y)} \mathbb{P}[X \neq Y], \tag{2.4}$$

where the infimum is over all couplings (X,Y) of (μ,ν). We recall a simple result bounding the total variation of product chains (see e.g. pg 59 of Levin, Peres, and Wilmer (2008)): for $n \in \mathbb{N}$ and $1 \leq i \leq n$, let μ_i and ν_i be measures on a finite space Ω_i and define measures μ and ν on $\prod_{i=1}^n \Omega_i$ by $\mu := \prod_{i=1}^n \mu_i$ and $\nu := \prod_{i-1}^n \nu_i$. Then

$$\|\mu - \nu\|_{\mathrm{TV}} \leq \sum_{i=1}^n \|\mu_i - \nu_i\|_{\mathrm{TV}}. \tag{2.5}$$

Recall equation (1.1) as the definition of the mixing time of a continuous-time Markov process. We will require several general mixing-time bounds throughout this work, which we present here.

Proposition 2.2 (Levin et al. (2008))**.** *Let X be a Markov process on a finite state space. Then for every $\varepsilon_1, \varepsilon_2 \in (0, 1/2)$,*

$$T_X(\varepsilon_2) \leq \left\lceil \frac{\log \varepsilon_2}{\log(2\varepsilon_1)} \right\rceil T_X(\varepsilon_1).$$

Proposition 2.3. *For any $m, n \in \mathbb{N}$,*

$$T_{\mathrm{RW}(2^m, f, G)}(2^{-n}) \leq (n+m) T_{\mathrm{RW}(f,G)}(1/4).$$

Proof. This follows by combining Proposition 2.2 with (2.5). □

Proposition 2.4 (Aldous and Fill (2002))**.** *Let X be a Markov process on a finite state space Ω with symmetric transition rates. Then the equilibrium distribution is uniform over Ω and for all $0 < \varepsilon < 1/2$ and $t \geq 2T_X(\varepsilon)$,*

$$\mathbb{P}(X_t = \omega_2 \,|\, X_0 = \omega_1) \geq \frac{(1-2\varepsilon)^2}{|\Omega|},$$

for all $\omega_1, \omega_2 \in \Omega$.

2.4 Failure of negative correlation

We conclude this preliminary section with a quick example to demonstrate that the exclusion process on a hypergraph does not enjoy the negative correlation property satisfied by the exclusion process on a graph. We first recall the version of the negative correlation property of the exclusion process on a graph to which we refer. The proof may be found in Liggett (1985): let $A \subset V$ and let $(\mathbf{x}_t)_{t \geq 0}$ be a 2-particle exclusion process on a graph $G = (V, E)$. Then for every $t \geq 0$,

$$\mathbb{P}[\mathbf{x}_t(1) \in A,\, \mathbf{x}_t(2) \in A] \leq \mathbb{P}[\mathbf{x}_t(1) \in A]\mathbb{P}[\mathbf{x}_t(2) \in A].$$

Now suppose $G = (V, E)$ is the hypergraph with $V = E = \{1, 2, 3, 4\}$ (i.e. there is only one edge), and that f is concentrated uniformly on the six possible 4-cycles. Suppose $(u_t^{\mathrm{RW}})_{t \geq 0}$ and $(v_t^{\mathrm{RW}})_{t \geq 0}$ are two independent realisations of $\mathrm{RW}(f, G)$, and $(\mathbf{x}_t^{\mathrm{EX}})_{t \geq 0}$ is a realisation of $\mathrm{EX}(2, f, G)$, with $\{u_0^{\mathrm{RW}}, v_0^{\mathrm{RW}}\} = \mathbf{x}_0^{\mathrm{EX}} = \{1, 2\}$. Let $A = \{3, 4\}$. We claim that there exist values of t such that

$$\mathbb{P}[u_t^{\mathrm{RW}} \in A, v_t^{\mathrm{RW}} \in A] < \mathbb{P}[\mathbf{x}_t^{\mathrm{EX}} \in A]. \tag{2.6}$$

Indeed, since the event $\{u_t^{\mathrm{RW}} \in A\}$ is less likely than seeing at least one incident in a unit-rate Poisson process by time t, we have

$$\mathbb{P}[u_t^{\mathrm{RW}} \in A, v_t^{\mathrm{RW}} \in A] = \mathbb{P}[u_t^{\mathrm{RW}} \in A]\mathbb{P}[v_t^{\mathrm{RW}} \in A] \leq (1 - e^{-t})^2.$$

On the other hand, the event $\{\mathbf{x}_t^{\mathrm{EX}} \in A\}$ is at least as likely as the edge ringing exactly once by time t, with the chosen permutation satisfying $\sigma(\{1, 2\}) = A$. That is,

$$\mathbb{P}[\mathbf{x}_t^{\mathrm{EX}} \in A] \geq \frac{1}{3} t e^{-t}.$$

Inequality (2.6) is therefore satisfied for any $t < 0.33$.

3 From 4-particle interchange to 2-particle exclusion

The aim of this section is to prove Lemma 1.9. We begin by characterizing every connected hypergraph in terms of how long it takes two independent random walkers on the hypergraph to arrive onto the same edge, which then rings for one of the walkers – a time we shall refer to as the *meeting time* of the two walkers (note that we do not require the two walkers to actually occupy the same vertex). It will be useful to consider such times, as we will be able to couple two independent walkers with a 2-particle interchange process, up until this meeting time (see Proposition 3.9 for this statement).

Formalising this, for $\mathbf{y} \in V^2$, let $(\mathbf{y}_t^{\mathrm{RW}})_{t\geq 0}$ be a realisation of $\mathrm{RW}(2, f, G)$ with $\mathbf{y}_0^{\mathrm{RW}} = \mathbf{y}$. Denote by Λ^1 and Λ^2 the Poisson processes used to generate the edge-ringing times of the two realisations, and let $\{e_n^1\}_{n\in\mathbb{N}}$ and $\{e_n^2\}_{n\in\mathbb{N}}$ be the two sequences of edge-choices (all as in Section 2.2). Define $M^{\mathrm{RW}}(\mathbf{y})$ to be the first time $\mathbf{y}_t^{\mathrm{RW}}(1)$ and $\mathbf{y}_t^{\mathrm{RW}}(2)$ are in the same edge which then rings in one of the processes:

$$M^{\mathrm{RW}}(\mathbf{y}) := \inf\left\{t > 0 : \exists e \in \{e^1_{\Lambda^1[0,t]}, e^2_{\Lambda^2[0,t]}\} \text{ with } \mathbf{y}^{\mathrm{RW}}_{t-}(1), \mathbf{y}^{\mathrm{RW}}_{t-}(2) \in e\right\}.$$

Definition 3.1. We say that a hypergraph G (satisfying Assumption 1.1) is *easy* if

$$\sup_{\mathbf{y}\in V^2} \mathbb{P}[M^{\mathrm{RW}}(\mathbf{y}) > 10^{10} T_{\mathrm{EX}(2,f,G)}(1/4)] \leq 1/1000.$$

Remark 3.2. We note that this definition is similar to Definition 4.1 of Oliveira (2013). One difference in the case of hypergraphs is that at the meeting time we cannot guarantee that the two independent walkers occupy the same site. This results in the analysis being slightly more challenging. For dealing with the case of easy hypergraphs, our strategy is to couple two realisations of $\mathrm{IP}(4, f, G)$ by, one by one, matching up the locations of the four particles under each process. An issue that arises when trying to do this (which is not a problem for standard graphs) is that it may happen that creating a new match will simultaneously destroy an existing match. Part 1 of Assumption 1.1 (that each edge contains at least seven vertices) plays a crucial role here in controlling these bad events.

Remark 3.3. Recall the d-uniform complete hypergraph from Example 1 in Section 1.2. There we saw that the mixing time of $\mathrm{EX}(2, f, G)$ satisfies

$$\frac{c_1}{\binom{|V|-1}{d-1}} \leq T_{\mathrm{EX}(2,f,G)}(1/4) \leq \frac{c_2}{\binom{|V|-1}{d-1}}$$

for some constants $c_1 \leq c_2$ (not depending on $|V|$). For this hypergraph, $M^{\mathrm{RW}}(\mathbf{y})$ has an exponential distribution with rate $\binom{|V|-2}{d-2}$ (the number of edges containing both $\mathbf{y}(1)$ and $\mathbf{y}(2)$). Thus

$$\mathbb{P}\left[M^{\mathrm{RW}}(\mathbf{y}) > \frac{10^{10}c_2}{\binom{|V|-1}{d-1}}\right] \leq \mathbb{P}[M^{\mathrm{RW}}(\mathbf{y}) > 10^{10} T_{\mathrm{EX}(2,f,G)}(1/4)] \leq \mathbb{P}\left[M^{\mathrm{RW}}(\mathbf{y}) > \frac{10^{10}c_1}{\binom{|V|-1}{d-1}}\right],$$

i.e.

$$\exp\left(-10^{10}c_2\frac{(d-1)}{(|V|-1)}\right) \leq \mathbb{P}[M^{\mathrm{RW}}(\mathbf{y}) > 10^{10} T_{\mathrm{EX}(2,f,G)}(1/4)] \leq \exp\left(-10^{10}c_1\frac{(d-1)}{(|V|-1)}\right).$$

With d fixed, we see that the d-uniform complete hypergraph is easy for small values of $|V|$, but non-easy when $|V|$ is sufficiently large.

Throughout the remainder of Section 3 we set

$$T := T_{\mathrm{EX}(2,f,G)}(1/4).$$

3.1 From 4-particle interchange to 2-particle exclusion: easy hypergraphs

We shall describe a coupling of two realisations of $\mathrm{IP}(4, f, G)$ in such a way that each pair of particles with the same label eventually occupies a common vertex (and then moves together); this is inspired by the coupling of two interchange processes in Chapter 14 of Aldous and Fill (2002).

For $\mathbf{x}, \mathbf{y} \in (V)_4$, let $(\mathbf{x}_t^{\mathrm{IP}})_{t\geq 0}$ and $(\mathbf{y}_t^{\mathrm{IP}})_{t\geq 0}$ be realisations of $\mathrm{IP}(4, f, G)$, started from $\mathbf{x}$ and $\mathbf{y}$ respectively, and constructed as follows. For each $1 \leq i \leq 4$, we shall wait for times t at which we have $\mathbf{x}_t^{\mathrm{IP}}(i)$ and $\mathbf{y}_t^{\mathrm{IP}}(i)$ in a common edge. Once they are in a common edge, we couple the ringing times of this edge so that it rings at the same time for both processes. If this edge rings before any of the other edges containing $\mathbf{x}_t^{\mathrm{IP}}(i)$ and also before any of the other edges containing $\mathbf{y}_t^{\mathrm{IP}}(i)$, then we get an opportunity to *match up* the two corresponding particles (specifically, to force them to occupy the same vertex). We will do this in a way that doesn't break any already matched pairs, so that the number of matches is non-decreasing over time. Our aim is to upper-bound how long it takes for all four pairs to become matched. To be precise, we say that the particle initially at $\mathbf{x}(i)$ and particle initially at $\mathbf{y}(i)$ (in the other interchange process) are *matched at time* t if $\mathbf{x}_t^{\mathrm{IP}}(i) = \mathbf{y}_t^{\mathrm{IP}}(i)$. We shall consider four intervals of time of equal length κT for some positive constant κ to be chosen later. Our strategy is to match up the ith pair of particles during time interval $[(i-1)\kappa T, i\kappa T)$ with a probability we can bound from below: the proof involves a relatively straightforward coupling argument; details can be found in Appendix A.

Lemma 3.4. *Fix $\delta > 0$ and $1 \leq i \leq 4$. Then there exists $\kappa > 0$ such that for any easy hypergraph G with probability at least $1-\delta$ we can match up the ith pair during $[(i-1)\kappa T, i\kappa T)$ without breaking any of the $i-1$ existing matches. Moreover, we can take $\kappa = 10^{12}\lceil \log(\delta^{-1})\rceil$.*

The proof of Lemma 1.9 for easy hypergraphs now follows trivially from the following result.

Lemma 3.5. *For any easy hypergraph G with probability at least 3/4 we can match up all 4 pairs by time $10^{14}T$.*

Proof. By Lemma 3.4, for any δ we can find a κ such that by time $4\kappa T$, we can match up all four pairs with probability at least $1-\delta$. Hence for any $\mathbf{x}, \mathbf{y} \in (V)_4$, the union bound gives

$$\mathbb{P}[\mathbf{x}_{4\kappa T}^{\mathrm{IP}} \neq \mathbf{y}_{4\kappa T}^{\mathrm{IP}}] \leq 4\delta.$$

We obtain the desired result by choosing $\delta = 1/16$ which gives $4\kappa = 12 \times 10^{12}$. □

Remark 3.6. Suppose that the following condition (which is stronger than that satisfied by an easy hypergraph) holds:

$$\sup_{\mathbf{y}\in V^2} \mathbb{P}[M^{\mathrm{RW}}(\mathbf{y}) > 10^{10} T_{\mathrm{RW}(f,G)}(1/4)] \leq 1/1000. \tag{3.1}$$

Since the proofs of Lemmas 3.4 and 3.5 don't depend on the value of T, in this case we are able to deduce that the mixing time of $\mathrm{IP}(4, f, G)$ is the same (up to multiplication by some universal constant) as the mixing time of $\mathrm{RW}(f, G)$. A similar condition to (3.1) is used as the definition of an easy graph in Oliveira (2013). However, it would appear that the negative correlation property (which, as we have seen, is enjoyed by the interchange process on graphs, but *not* on hypergraphs) is required in order to be able to deal with G not satisfying (3.1). As such, we are forced to define easy hypergraphs in terms of $T_{\mathrm{EX}(2,f,G)}$, leading to a slightly weaker conclusion in Lemma 1.9 than we might like.

3.2 From 4-particle interchange to 2-particle exclusion: non-easy hypergraphs

We begin with a result showing that for non-easy hypergraphs the average meeting time for two independent random walkers is unlikely to be quick. Intuitively, this follows from the following observations. We know there exists a pair of vertices such that random walkers started from these two states likely take a long time to meet. If we look at where these two walkers are at time of order $T_{\mathrm{RW}(f,G)}(1/4)$, they will be close to uniform. Hence, starting random walkers from a uniform pair we see that they will likely still take a long time to meet. The proofs of Lemmas 3.7 and 3.8, and of Proposition 3.9 are (somewhat technical) extensions of corresponding results of Oliveira (2013), and can be found in Appendix A.

Lemma 3.7. *For every non-easy hypergraph satisfying Assumption 1.1, we have*

$$\sum_{\mathbf{v}\in V^2} \frac{\mathbb{P}\big[M^{\mathrm{RW}}(\mathbf{v}) \le 20T\big]}{|V|^2} \le \frac{1}{1000}.$$

For $\mathbf{x} \in V^4$, let $\mathbf{x}_t^{\mathrm{RW}}$ be a realisation of $\mathrm{RW}(4, f, G)$ with $\mathbf{x}_0^{\mathrm{RW}} = \mathbf{x}$. For such an $\mathbf{x}$, define $\bar{M}^{\mathrm{RW}}(\mathbf{x})$ to be the first time *any two* of $\mathbf{x}_t^{\mathrm{RW}}(1)$, $\mathbf{x}_t^{\mathrm{RW}}(2)$, $\mathbf{x}_t^{\mathrm{RW}}(3)$, $\mathbf{x}_t^{\mathrm{RW}}(4)$ first arrive onto the same edge which then rings for one of them. Formally,

$$\bar{M}^{\mathrm{RW}}(\mathbf{x}) := \inf\Big\{t : 0 : \exists\, 1 \le i < j \le 4, \exists\, e \in \{e^i_{\mathcal{P}^i[0,t]}, e^j_{\mathcal{P}^j[0,t]}\} \text{ with } \mathbf{x}^{\mathrm{RW}}_{t-}(i) \in e,\, \mathbf{x}^{\mathrm{RW}}_{t-}(j) \in e\Big\}.$$

Lemma 3.8. *Let $\mathbf{x} \in (V)_4$. Then for any $\varepsilon \in (0, 1)$,*

$$\mathbb{P}[\bar{M}^{\mathrm{RW}}(\mathbf{x}^{\mathrm{IP}}_{20T}) \le 20T] \le 12(\varepsilon + \varepsilon^{-1}2^{-20}) + 25(1+\varepsilon) \sum_{\mathbf{u}\in V^2} \frac{\mathbb{P}[M^{\mathrm{RW}}(\mathbf{u}) \le 20T]}{|V|^2}$$

Next, we provide a bound which relates the total-variation distance between two interchange processes with four particles to the probability that any two of four independent walkers have 'met'.

Proposition 3.9. *For any $\mathbf{x} \in (V)_4$ and $s \ge 0$:*

$$\|\mathcal{L}[\mathbf{x}^{\mathrm{RW}}_s] - \mathcal{L}[\mathbf{x}^{\mathrm{IP}}_s]\|_{\mathrm{TV}} \le \mathbb{P}[\bar{M}^{\mathrm{RW}}(\mathbf{x}) \le s].$$

Lemma 3.10. *For every non-easy hypergraph G satisfying Assumption 1.1 and any two realisations of* $\mathrm{IP}(4, f, G)$*, denoted* $\{\mathbf{x}_t^{\mathrm{IP}}\}$ *and* $\{\mathbf{y}_t^{\mathrm{IP}}\}$*, we have*

$$\|\mathcal{L}[\mathbf{x}_{40T}^{\mathrm{IP}}] - \mathcal{L}[\mathbf{y}_{40T}^{\mathrm{IP}}]\|_{\mathrm{TV}} \leq \mathbb{P}[\bar{M}^{\mathrm{RW}}(\mathbf{x}_{20T}^{\mathrm{IP}})) \leq 20T] + \mathbb{P}[\bar{M}^{\mathrm{RW}}(\mathbf{y}_{20T}^{\mathrm{IP}}) \leq 20T] + 2^{-18}.$$

Proof. By Proposition 3.9 and the triangle inequality for total-variation, for any $\mathbf{u}, \mathbf{v} \in (V)_4$,

$$\|\mathcal{L}[\mathbf{u}_{20T}^{\mathrm{IP}}] - \mathcal{L}[\mathbf{v}_{20T}^{\mathrm{IP}}]\|_{\mathrm{TV}} \leq \mathbb{P}[\bar{M}^{\mathrm{RW}}(\mathbf{u}) \leq 20T] + \mathbb{P}[\bar{M}^{\mathrm{RW}}(\mathbf{v}) \leq 20T] + \|\mathcal{L}[\mathbf{u}_{20T}^{\mathrm{RW}}] - \mathcal{L}[\mathbf{v}_{20T}^{\mathrm{RW}}]\|_{\mathrm{TV}}. \tag{3.2}$$

For any $A \subset (V)_4$,

$$\begin{aligned}\mathbb{P}[\mathbf{x}_{40T}^{\mathrm{IP}} \in A] - \mathbb{P}[\mathbf{y}_{40T}^{\mathrm{IP}} \in A] &= \mathbb{E}\big[\mathbb{P}[\mathbf{x}_{40T}^{\mathrm{IP}} \in A |\, \mathbf{x}_{20T}^{\mathrm{IP}}] - \mathbb{P}[\mathbf{y}_{40T}^{\mathrm{IP}} \in A |\, \mathbf{y}_{20T}^{\mathrm{IP}}]\big] \\ &\leq \mathbb{E}\big[\|\mathbb{P}[\mathbf{x}_{40T}^{\mathrm{IP}} \in \cdot |\, \mathbf{x}_{20T}^{\mathrm{IP}}] - \mathbb{P}[\mathbf{y}_{40T}^{\mathrm{IP}} \in \cdot |\, \mathbf{y}_{20T}^{\mathrm{IP}}]\|_{\mathrm{TV}}\big].\end{aligned}$$

Maximizing over A and using the Markov property gives

$$\|\mathcal{L}[\mathbf{x}_{40T}^{\mathrm{IP}}] - \mathcal{L}[\mathbf{y}_{40T}^{\mathrm{IP}}]\|_{\mathrm{TV}} \leq \mathbb{E}\big[\|\mathbb{P}[(\mathbf{x}_{20T}^{\mathrm{IP}})_{20T}^{\mathrm{IP}} \in \cdot |\, \mathbf{x}_{20T}^{\mathrm{IP}}] - \mathbb{P}[(\mathbf{y}_{20T}^{\mathrm{IP}})_{20T}^{\mathrm{IP}} \in \cdot |\, \mathbf{y}_{20T}^{\mathrm{IP}}]\|_{\mathrm{TV}}\big].$$

We now make use of equation (3.2) with $\mathbf{u} = \mathbf{x}_{20T}^{\mathrm{IP}}$ and $\mathbf{v} = \mathbf{y}_{20T}^{\mathrm{IP}}$ to give

$$\begin{aligned}\|\mathcal{L}[\mathbf{x}_{40T}^{\mathrm{IP}}] - \mathcal{L}[\mathbf{y}_{40T}^{\mathrm{IP}}]\|_{\mathrm{TV}} \leq \mathbb{E}\big[\mathbb{P}[\bar{M}^{\mathrm{RW}}(\mathbf{x}_{20T}^{\mathrm{IP}}) \leq 20T |\, \mathbf{x}_{20T}^{\mathrm{IP}}] + \mathbb{P}[\bar{M}^{\mathrm{RW}}(\mathbf{y}_{20T}^{\mathrm{IP}}) \leq 20T |\, \mathbf{y}_{20T}^{\mathrm{IP}}]\big] \\ + \sup_{\mathbf{u},\mathbf{v} \in (V)_4} \|\mathcal{L}[\mathbf{u}_{20T}^{\mathrm{RW}}] - \mathcal{L}[\mathbf{v}_{20T}^{\mathrm{RW}}]\|_{\mathrm{TV}}.\end{aligned}$$

Using Proposition 2.3 for the third term on the right-hand side, gives the desired result. □

We are now ready to prove the main result of this subsection.

Proof of Lemma 1.9 for non-easy hypergraphs. We in fact show that for any two realisations of $\mathrm{IP}(4, f, G)$, denoted $\{\mathbf{x}_t^{\mathrm{IP}}\}$ and $\{\mathbf{y}_t^{\mathrm{IP}}\}$, we have

$$\|\mathcal{L}[\mathbf{x}_{40T}^{\mathrm{IP}}] - \mathcal{L}[\mathbf{y}_{40T}^{\mathrm{IP}}]\|_{\mathrm{TV}} \leq 1/4.$$

Combining Lemmas 3.8 and 3.10 we have that for every $\varepsilon \in (0, 1)$,

$$\|\mathcal{L}[\mathbf{x}_{40T}^{\mathrm{IP}}] - \mathcal{L}[\mathbf{y}_{40T}^{\mathrm{IP}}]\|_{\mathrm{TV}} \leq 24(\varepsilon + \varepsilon^{-1}2^{-19}) + 50(1+\varepsilon) \sum_{\mathbf{u} \in V^2} \frac{\mathbb{P}[M^{\mathrm{RW}}(\mathbf{u}) \leq 20T]}{|V|^2} + 2^{-18}.$$

Now by Lemma 3.7, this becomes

$$\|\mathcal{L}[\mathbf{x}_{40T}^{\mathrm{IP}}] - \mathcal{L}[\mathbf{y}_{40T}^{\mathrm{IP}}]\|_{\mathrm{TV}} \leq 24(\varepsilon + \varepsilon^{-1}2^{-19}) + \frac{50(1+\varepsilon)}{1000} + 2^{-18}.$$

Setting $\varepsilon = 10^{-3}$ completes the proof.

□

4 The chameleon process

The chameleon process was originally constructed by Morris (2006), and then adapted by Oliveira (2013) to analyse the mixing time of the interchange process on graphs (as opposed to hypergraphs). It is built on top of an underlying interchange process, with the aim of helping to describe the distribution of the location of the kth particle in this process, conditional on the locations of the $k-1$ other particles.

In the chameleon process there are always $|V|$ particles (as opposed to the k in the corresponding interchange), although not all particles are distinguishable from each other. In addition, each particle has an associated *colour*: one of black, red, pink and white. The movement of particles in the chameleon process follows that of the underlying interchange process, in the sense that the locations of particles in both processes are updated using the same functions I as described in the graphical construction of Section 2.2. The updates of the chameleon process which correspond to updates of the interchange may or may not (depending on the time of the update and the configuration of colours at this time) result in a recolouring of some of the particles. In the chameleon process of Oliveira (2013), if the vertices in a ringing edge contain a red and a white particle then both of these particles are recoloured pink. In the lazy version of the interchange process on a graph (in which nothing happens with probability 1/2 when an edge rings), when an edge rings with endpoints containing a red and a white particle, with probability 1/2 they switch places and with probability 1/2 they do not. Colouring both particles pink encodes the fact that at either vertex just after the edge rings we may have a white particle or a red particle, and that these are equally likely.

We wish to use this notion of *pinkening* to encode similar events in the interchange process on hypergraphs, but the situation here is quite different since more than two particles are moved when an edge rings, and the way in which they move depends on the permutation chosen. As a result, describing precisely when these pinkenings occur for our version of the chameleon process is rather complicated, but the underlying motivation can be explained relatively simply. Suppose that an edge e rings and a permutation σ is chosen to move the particles on that edge. To decide which particles to pinken, we make use of the fact that f_e is constant on conjugacy classes to construct another permutation $\tilde{\sigma}$ with the same cycle structure as σ. Figure 5 gives an indication of how $\tilde{\sigma}$ will be produced from knowledge of σ and the particle colours: by modifying the trajectories of four particular particles we are able to ensure that not only does $\tilde{\sigma}$ have the same cycle structure as σ, but that the trajectories of all black particles in the edge are identical under both permutations. We then look for vertices v such that under σ a red particle is moved to v and under $\tilde{\sigma}$ a white particle is moved to v; a certain subset of these particles will be pinkened. It is for this reason that we are able to relate the mixing time of k particles to that of just four particles in Lemma 1.8.

The chameleon process updates at additional times (compared to its corresponding interchange). We refer to these additional updates as *depinkings*, as at these times we get the opportunity to collectively recolour all pink particles in the system either red or white. As in (Oliveira, 2013), we will only perform a depinking once there are a large number of pink particles (compared to the number of red and white) in the system. In fact describing formally when a depinking occurs is very simple as it depends just on the time and on the numbers of particles of various colours.

In Section 4.1 we present the key property that we wish our chameleon process for hypergraphs

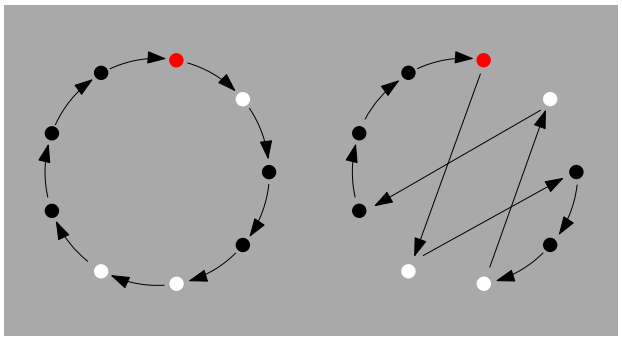

Figure 5: If σ is the cycle given by the arrows in the image on the left, then one possible candidate for $\tilde{\sigma}$ is the cycle given by the arrows in the image on the right. Note that all black particles follow the same trajectories under both cycles, and that the colour of $\tilde{\sigma}(v)$ agrees with that of $\sigma(v)$ for all v.

to possess, and show in Lemma 4.2 how this will help us to understand the distribution of particles in the interchange process. The rest of Section 4 is devoted to careful construction of a chameleon process satisfying this key property.

4.1 Connection between interchange and chameleon

Given a $(k-1)$-tuple $\mathbf{z} \in (V)_{k-1}$, let $\mathbf{O}(\mathbf{z}) := \{\mathbf{z}(1), \ldots, \mathbf{z}(k-1)\}$ be the (unordered) set of coordinates of $\mathbf{z}$. The chameleon process is a continuous-time Markov process with state-space

$$\Omega_k(V) := \{(\mathbf{z}, R, P, W) : \mathbf{z} \in (V)_{k-1}, \text{ and sets } \mathbf{O}(\mathbf{z}), R, P, W \text{ partition } V\}.$$

A particle at vertex v is said to be *red* if $v \in R$, *white* if $v \in W$, *pink* if $v \in P$ and *black* if $v \in \mathbf{O}(\mathbf{z})$. Note that, due to the nature of the state-space, we can distinguish between the various black particles, whereas two red/white/pink particles are indistinguishable from each other.

We denote the state at time t of the chameleon process started from $M_0 = (\mathbf{z}, R, P, W)$ as

$$M_t = (\mathbf{z}_t, R_t, P_t, W_t).$$

We say that a particle at vertex v at time t is *red at time* t if $v \in R_t$ (and similarly for the other colours). We define now a notion of *ink*, which represents the amount of *redness* either at a vertex or in the whole system. A vertex v has 1 unit of ink at time t if $v \in R_t$ and half a unit if $v \in P_t$. Formally then, we define for each $v \in V$ and $t \geq 0$,

$$\mathrm{ink}_t(v) := \mathbf{1}_{\{v \in R_t\}} + \frac{1}{2}\mathbf{1}_{\{v \in P_t\}}. \tag{4.1}$$

In order for the chameleon process to aid us in understanding the conditional distribution of the kth particle in an interchange process with k particles (conditioned on the trajectory of the other $k-1$ particles), we wish for the chameleon process to satisfy several properties – these are formulated in terms of the first update time of the chameleon process. The behaviour of the chameleon process is not time-homogeneous; as a result we state the desired properties in terms of the first jump time after some fixed time s. For every $s \geq 0$ and $n \in \mathbb{N}$, we denote

by τ_n^s the time we have to wait after time s for the nth incident (after time s) of the Poisson process underlying the interchange process. Similarly, we denote by $\bar{\tau}_n^s$ the time we have to wait after time s for the nth update of the chameleon process. We write τ_n and $\bar{\tau}_n$ for τ_n^0 and $\bar{\tau}_n^0$, respectively. For an interchange process $\{\mathbf{x}_t^{\mathrm{IP}}\}_{t\geq s}$, we denote by σ_n^s the permutation chosen to be applied at time $s+\tau_n^s$.

Box 4.1. Desired properties of chameleon process
For any $s \geq 0$, $\mathbf{z} \in (V)_{k-1}$ and $M_0 = (\mathbf{z}, R, P, W) \in \Omega_k(V)$, and any chameleon process $\{M_t\}_{t\geq s}$ with $M_s := M_0$, defined with respect to some interchange process $\{\mathbf{x}_t^{\mathrm{IP}}\}_{t\geq s}$ with $\mathbf{x}_s^{\mathrm{IP}} := \mathbf{x} = (\mathbf{z}, x) \in (V)_k$:

1. if $\bar{\tau}_1^s = \tau_1^s$ (i.e. the first update of M corresponds to an incident in $\mathbf{x}^{\mathrm{IP}}$), then all of the following hold:
   (a) there exists a permutation $\tilde{\sigma}_1^s$ such that $\sigma_1^s(\mathbf{z}) = \tilde{\sigma}_1^s(\mathbf{z})$ almost surely;
   (b) setting $\sigma_n^s = \tilde{\sigma}_n^s$ for all $n > 1$, the process $\{\tilde{\mathbf{x}}_t^{\mathrm{IP}}\}_{t\geq s}$ with $\tilde{\mathbf{x}}_s^{\mathrm{IP}} = \mathbf{x}$, which uses the same Poisson process of jump times and same edge choices as $\mathbf{x}^{\mathrm{IP}}$, but permutation choices $\tilde{\sigma}_n$, has the same transition kernel as $\mathbf{x}^{\mathrm{IP}}$;
   (c) almost surely, for all $v \in V$, the ink process under M satisfies
   $$\mathrm{ink}_{s+\tau_1^s}(v) = \frac{1}{2}\mathrm{ink}_s((\sigma_1^s)^{-1}(v)) + \frac{1}{2}\mathrm{ink}_s((\tilde{\sigma}_1^s)^{-1}(v)).$$
2. if $\bar{\tau}_1^s < \tau_1^s$ (i.e. a depinking at time $\bar{\tau}_1^s$), then for all $v \in V$, $\mathbb{E}[\mathrm{ink}_{\bar{\tau}_1^s}(v)] = \mathrm{ink}_0(v)$.

The following result, which is based upon Lemma 1 of Morris (2006), shows how the distribution of particles in the interchange process on a hypergraph is related to the evolution of ink in the corresponding chameleon process.

Lemma 4.2. *Suppose the properties in Box 4.1 are satisfied. Fix $\mathbf{x} = (\mathbf{z}, x) \in (V)_k$ and $M_0 = (\mathbf{z}, \{x\}, \varnothing, V \setminus (\mathbf{O}(\mathbf{z}) \cup \{x\})) \in \Omega_k(V)$. (That is, M_0 contains one red particle at vertex x.) Let $\{M_t\}_{t\geq 0}$ be a chameleon process started from M_0 with corresponding interchange process $\{\mathbf{x}_t^{\mathrm{IP}}\}_{t\geq 0}$.*

Then for every $t > 0$, and $\mathbf{b} = (\mathbf{c}, b) \in (V)_k$,

$$\mathbb{P}[\mathbf{x}_t^{\mathrm{IP}} = \mathbf{b}] = \mathbb{E}[\mathrm{ink}_t(b)\mathbf{1}_{\{\mathbf{z}_t = \mathbf{c}\}}].$$

Proof. Suppose $\{\mathbf{x}_t^{\mathrm{IP}}\}$ evolves using permutations $\{\sigma_n\}_{n\in\mathbb{N}}$ (see Section 2.2).

For each $j \in \mathbb{N}$, consider a process $\{M_t^j\}_{t\geq 0}$ which is identical to $\{M_t\}_{t\geq 0}$ for all $t < \bar{\tau}_j$ but evolves as the interchange process (i.e. with no further recolourings) for all $t \geq \bar{\tau}_j$. More formally, for all $t \geq \bar{\tau}_j$,

$$M_t^j = (I_{(\bar{\tau}_j, t]}(\mathbf{z}_{\bar{\tau}_j}), I_{(\bar{\tau}_j, t]}(R_{\bar{\tau}_j}), I_{(\bar{\tau}_j, t]}(P_{\bar{\tau}_j}), I_{(\bar{\tau}_j, t]}(W_{\bar{\tau}_j})),$$

where I is the map used in the graphical construction (of Section 2.2) of the interchange process $\{\mathbf{x}_t^{\mathrm{IP}}\}$.

Notice that the almost-sure limit of $\{M_t^j\}_{t\geq 0}$ as $j \to \infty$ is the chameleon process $\{M_t\}_{t\geq 0}$. As a result, by the dominated convergence theorem, it suffices to prove that for each $j \in \mathbb{N}$ and $b \in V$,

$$\mathbb{P}[x_t^{\mathrm{IP}} = b |\, \mathbf{z}_t^{\mathrm{IP}}] = \mathbb{E}[\mathrm{ink}_t^j(b) |\, \mathbf{z}_t^{\mathrm{IP}}],$$

where $\mathrm{ink}_t^j(b)$ is the amount of ink under M_t^j at vertex b. We prove this by induction on j. The case $j = 1$ is trivial since the particle initially at x is the only red particle (and there are no pink particles). For the inductive step we wish to show that

$$\mathbb{E}[\mathrm{ink}_t^j(b) |\, \mathbf{z}_t^{\mathrm{IP}}] = \mathbb{E}[\mathrm{ink}_t^{j+1}(b) |\, \mathbf{z}_t^{\mathrm{IP}}].$$

For $t < \bar{\tau}_j$, these are equal since the two processes evolve identically for such times. The update at time $\bar{\tau}_j$ of process $\{M_t^{j+1}\}$ is a chameleon step and could be of two types: also an update of the interchange process (i.e. $\bar{\tau}_j$ is a point in the Poisson process of update times of the interchange), or not. Suppose we are in the first case. We condition on the common state of M^j and M^{j+1} at time $\bar{\tau}_{j-1}$. We want to show that almost surely

$$\mathbb{E}\left[\mathbb{E}[\mathrm{ink}_{\bar{\tau}_j}^j(b) |\, \mathbf{z}_{\bar{\tau}_j}^{\mathrm{IP}}, M_{\bar{\tau}_{j-1}}^j]\right] = \mathbb{E}\left[\mathbb{E}[\mathrm{ink}_{\bar{\tau}_j}^{j+1}(b) |\, \mathbf{z}_{\bar{\tau}_j}^{\mathrm{IP}}, M_{\bar{\tau}_{j-1}}^{j+1}]\right]. \tag{4.2}$$

By the strong Markov property at time $\bar{\tau}_{j-1}$, and since the desired properties of Box 4.1 hold, we can construct a chameleon process $\{\tilde{M}_t^j\}$ (with associated interchange process $\tilde{\mathbf{x}}^{\mathrm{IP}}$) which is identical to $\{M_t\}$ for all $t < \bar{\tau}_j$ but evolves as an interchange process (i.e. with no further recolourings) using permutation choices $\{\tilde{\sigma}_n\}$ for all $t \geq \bar{\tau}_j$. By property 1c of Box 4.1

$$\frac{1}{2}\mathrm{ink}_{\bar{\tau}_j}^j(b) + \frac{1}{2}\widetilde{\mathrm{ink}}_{\bar{\tau}_j}^j(b) = \mathrm{ink}_{\bar{\tau}_j}^{j+1}(b)\,,$$

for all $b \in V$, almost surely (where $\widetilde{\mathrm{ink}}$ is the ink process under $\tilde{M}^j$). Since $\{M_t^j\}$ and $\{\tilde{M}_t^j\}$ have the same distribution (thanks to property 1b) and the trajectory of $\mathbf{z}$ is identical under each process (property 1a),

$$\begin{aligned}\mathbb{E}\left[\mathbb{E}[\mathrm{ink}_{\bar{\tau}_j}^j(b) |\, \mathbf{z}_{\bar{\tau}_j}^{\mathrm{IP}}, M_{\bar{\tau}_{j-1}}^j]\right] &= \frac{1}{2}\mathbb{E}\left[\mathbb{E}[\mathrm{ink}_{\bar{\tau}_j}^j(b) |\, \mathbf{z}_{\bar{\tau}_j}^{\mathrm{IP}}, M_{\bar{\tau}_{j-1}}^j]\right] + \frac{1}{2}\mathbb{E}\left[\mathbb{E}[\widetilde{\mathrm{ink}}_{\bar{\tau}_j}^j(b) |\, \tilde{\mathbf{z}}_{\bar{\tau}_j}^{\mathrm{IP}}, \tilde{M}_{\bar{\tau}_{j-1}}^j]\right] \\ &= \mathbb{E}\left[\mathbb{E}[\mathrm{ink}_{\bar{\tau}_j}^{j+1}(b) |\, \mathbf{z}_{\bar{\tau}_j}^{\mathrm{IP}}, M_{\bar{\tau}_{j-1}}^{j+1}]\right].\end{aligned}$$

We are left to deal with the second case, when $\bar{\tau}_j$ is not an update of the interchange process. In this case there must be a depinking at time $\bar{\tau}_j$. We wish to show (4.2) holds, so again use the strong Markov property at time $\bar{\tau}_{j-1}$ to obtain

$$\mathbb{E}\left[\mathbb{E}[\mathrm{ink}_{\bar{\tau}_j}^{j+1}(b) |\, \mathbf{z}_{\bar{\tau}_j}^{\mathrm{IP}}, M_{\bar{\tau}_{j-1}}^{j+1}]\right] = \mathbb{E}\left[\mathbb{E}[\mathrm{ink}_{\bar{\tau}_j}^{j+1}(b) |\, \mathbf{z}_{\bar{\tau}_{j-1}}^{\mathrm{IP}}, M_{\bar{\tau}_{j-1}}^{j+1}]\right] = \mathbb{E}\left[\mathbb{E}[\mathrm{ink}_{\bar{\tau}_j}^{j}(b) |\, \mathbf{z}_{\bar{\tau}_{j-1}}^{\mathrm{IP}}, M_{\bar{\tau}_{j-1}}^{j+1}]\right],$$

where the second equality follows from property 2 of Box 4.1. This completes the proof. □

4.2 New permutations from old

Our goal now is to construct the chameleon process in such a way that the desired properties in Box 4.1 are satisfied. The first step to this is to detail how to produce the permutations $\{\tilde{\sigma}_n\}$ used in the other interchange process $\tilde{\mathbf{x}}^{\mathrm{IP}}$ described there. In this section we construct a function ρ which will be later used to construct this sequence of permutations. The function takes as its inputs a permutation σ and a subset A of the vertices, and outputs a permutation $\tilde{\sigma}$. Crucially, the function has the property that if the input permutation is uniform on cycles then also the output permutation is uniform on cycles. Furthermore if we average over the input permutation, the output permutation is independent of the choice of A.

Recall that every permutation can be written in canonical cyclic decomposition form: the function ρ will act independently on each cycle in a given permutation's decomposition, and so we begin by describing its action on a cycle.

To that end, for $d \geq 4$ let $[d] = \{1, 2, \ldots, d\}$, and let $d' = \lfloor d/4 \rfloor$ (the floor of $d/4$). We shall write $\mathcal{C}_d$ for the conjugacy class of $\mathcal{S}_d$ consisting of cycles of length d.

For $i = 1, \ldots, d'$ define the function $\rho_i : [d] \to [d]$ by

$$\rho_i(j) = \begin{cases} 2d' + 2i - 1 & j = 2i - 1 \\ 2i - 1 & j = 2d' + 2i - 1 \\ j & \text{otherwise.} \end{cases} \tag{4.3}$$

For a cycle $\sigma \in \mathcal{C}_d$ we may write

$$\sigma = (\sigma^0(1)\ \sigma(1)\ \sigma^2(1)\ \ldots\ \sigma^{d-1}(1))$$

where for $m \in \mathbb{N}$ we write σ^m for the composition of m copies of σ (and where we may sometimes write $\sigma^d \equiv \sigma^0$ for the identity permutation). Given a cycle σ, we define $\rho_i(\sigma) \in \mathcal{C}_d$ to be the permutation satisfying

$$\rho_i(\sigma)^j(1) = \sigma^{\rho_i(j)}(1)\,, \quad j = 1, \ldots, d.$$

This permutation is clearly a cycle, and may be written as

$$\rho_i(\sigma) = (\sigma^0(1)\ \ldots\ \sigma^{2i-2}(1)\ \sigma^{2d'+2i-1}(1)\ \sigma^{2i}(1)\ \ldots\ \sigma^{2d'+2i-2}(1)\ \sigma^{2i-1}(1)\ \sigma^{2d'+2i}(1)\ \ldots\ \sigma^{d-1}(1)).$$

Remark 4.3. Note that $\rho_i(\sigma)$ may be obtained from σ by multiplication (on the right) by the product of the two disjoint transpositions:

$$(\sigma^{2i-2}(1)\ \sigma^{2d'+2i-2}(1))(\sigma^{2i-1}(1)\ \sigma^{2d'+2i-1}(1))\,.$$

Lemma 4.4. *For any $i, j \in [d']$:*

1. *The function ρ_i is self-inverse;*
2. *functions ρ_i and ρ_j commute for $i \neq j$.*

Proof. Part 1 follows directly from the definition of ρ_i. Part 2 follows from the observation that the transpositions in Remark 4.3 corresponding to ρ_i and ρ_j commute for $i \neq j$. □

Definition 4.5. Given a set $A \subseteq [d']$ and a cycle $\sigma \in \mathcal{C}_d$, we define $\rho_A : [d] \to [d]$ to be the composition of the functions $\{\rho_i : i \in A\}$ appearing in (4.3). (Thanks to Part 2 of Lemma 4.4 this function is well-defined.) If $A = \varnothing$ then we set $\rho_A(\sigma) = \sigma$.
This in turn defines a cycle $\rho_A(\sigma) \in \mathcal{C}_d$ satisfying

$$\rho_A(\sigma)^j(1) = \sigma^{\rho_A(j)}(1)\,, \quad j = 1, \ldots, d. \tag{4.4}$$

We now state some facts about the permutation $\rho_A(\sigma)$ in the case that σ is a d-cycle, which all follow simply from the definition of ρ_i and Lemma 4.4. Write

$$H_i = \{2i-2, 2i-1, 2d'+2i-2, 2d'+2i-1\}\,,$$

and for $A \subseteq [d']$ let $H_A = \bigcup_{i \in A} H_i$.

Lemma 4.6. *For any $\sigma \in \mathcal{C}_d$ and set $A \subseteq [d']$:*

1. *$\rho_A(\sigma) \in \mathcal{C}_d$;*
2. *$\rho_A(\rho_A(\sigma)) = \sigma$;*
3. *for all $x \in [d]$, $\rho_A(\sigma)(x) = \sigma(x)$ unless $x = \sigma^j(1)$ for some $j \in H_A$.*

Now consider what happens when we apply the function ρ to a uniformly chosen cycle $\sigma \in \mathcal{C}_d$. Clearly, for any set A chosen independently of σ, the resulting cycle $\rho_A(\sigma)$ will also be uniformly distributed on $\mathcal{C}_d$. Most importantly, this remains true even when A is allowed to depend upon σ, as long as a certain condition is met, as explained in the following Lemma. We denote by $\mathcal{P}^\Omega$ the power set of a set Ω.

Lemma 4.7. *Suppose $A : \mathcal{C}_d \to \mathcal{P}^{[d']}$ satisfies for all $\sigma \in \mathcal{C}_d$,*

$$A(\rho_{A(\sigma)}(\sigma)) = A(\sigma)\,, \tag{4.5}$$

and that σ is chosen uniformly from $\mathcal{C}_d$. Then $\rho_{A(\sigma)}(\sigma)$ is also uniform on $\mathcal{C}_d$. Moreover, if we average over the input permutation σ, then the output $\rho_{A(\sigma)}(\sigma)$ is independent of the choice of A.

Proof. Given the cycle σ, let $\tilde{\sigma} = \rho_{A(\sigma)}(\sigma)$. The assumption on A says that $A(\tilde{\sigma}) = A(\sigma)$. It follows that

$$\rho_{A(\tilde{\sigma})}(\tilde{\sigma}) = \rho_{A(\sigma)}(\tilde{\sigma}) = \rho_{A(\sigma)}(\rho_{A(\sigma)}(\sigma)) = \sigma$$

by Part 2 of Lemma 4.6.

□

Although Lemma 4.7 is relatively simple, its importance should be emphasised at this point. We shall make use of the function ρ_A to generate the random permutations $\tilde{\sigma}_n$ used in the construction outlined in Box 4.1, and in doing so the input A will depend on the state of the chameleon process M. The second part of Lemma 4.7 will be used to guarantee that the permutation $\tilde{\sigma}_n = \rho_A(\sigma_n)$ is independent of A. (The permutations $\tilde{\sigma}$ will be used to generate

an interchange process $\tilde{\mathbf{x}}^{\mathrm{IP}}$, and so it will be crucial that these do not depend on the state of the process.)

In the rest of the paper we will want to apply the function ρ to general permutations $\sigma \in \mathcal{S}_e$ (i.e. not only cycles) for some $e \in E$, an edge of a graph. To accomplish this, we write σ in its canonical cyclic decomposition form:

$$\sigma = \sigma_1 \circ \cdots \circ \sigma_K,$$

where $K = K(\sigma)$ denotes the number of cycles in σ. Suppose the vertices V are labelled $\{1, \ldots, |V|\}$. Denote by $e_1, \ldots, e_K$ a partition of e according to the cycles of σ (that is, the set e_i is the non-trivial orbit of cycle σ_i), and for each $i \in \{1, \ldots, K\}$ set $d_i = |e_i|$, $d_i' = \lfloor d_i/4 \rfloor$, and $m_i = \min\{j \in e_i\}$. We also write $d = |e|$.

Now suppose that $A(\sigma)$ is a vector of the form

$$A(\sigma) = (A_1(\sigma_1), A_2(\sigma_2), \ldots, A_K(\sigma_K)), \tag{4.6}$$

where each $A_i : \mathcal{C}_{d_i} \to \mathcal{P}^{[d_i']}$. In this case we can easily extend the definition of ρ_A to apply to general permutations (not only cycles) as follows.

Definition 4.8. Let $\sigma \in \mathcal{S}_e$ be a permutation acting on the vertices of the edge e, and assume that A is a function on $\mathcal{S}_e$ satisfying (4.6). Then we define $\rho_{A(\sigma)}(\sigma)$ to be the product of the cycles obtained by applying the function $\rho_{A_i(\sigma_i)}$ to each cycle σ_i in the canonical decomposition of σ. Formally:

$$\rho_{A(\sigma)}(\sigma) = \prod_{i=1}^{K} \rho_{A_i(\sigma_i)}(\sigma_i),$$

where the cycle $\rho_{A_i(\sigma_i)}(\sigma_i)$ satisfies

$$\rho_{A_i(\sigma_i)}(\sigma_i)^j(m_i) = \sigma_i^{\rho_{A_i(\sigma_i)}(j)}(m_i).$$

Definition 4.8 is rather notationally cumbersome, but in essence it says that $\rho_A(\sigma)$ is obtained from σ by modifying each of its cycles independently, using functions ρ_{A_i} with which we are already familiar. We have observed above that $\rho_{A_i(\sigma_i)}(\sigma_i)$ is also a cycle of size d_i, and so it follows that $\rho_{A(\sigma)}(\sigma)$ will have the same cycle structure as σ. In addition, we see from Lemma 4.6 that the function ρ_A is self-inverse.

Furthermore, note that if we choose a random permutation $\sigma \in \mathcal{S}_e$ according to a law f_e on $\mathcal{S}_e$ which is constant on conjugacy classes, then given the sizes of the cycles in the decomposition of σ, the elements in each cycle are (marginally) uniform from e. We can therefore obtain a corollary to Lemma 4.7:

Corollary 4.9. *Suppose that A is a function on $\mathcal{S}_e$ satisfying (4.6), and that for all $\sigma \in \mathcal{S}_d$ and $i = 1, \ldots, K$,*

$$A_i(\rho_{A_i(\sigma_i)}(\sigma_i)) = A_i(\sigma_i).$$

Suppose in addition that σ is chosen according to law f_e on $\mathcal{S}_e$ for some f_e constant on conjugacy classes. Then $\rho_{A(\sigma)}(\sigma)$ also has law f_e on $\mathcal{S}_e$.

Proof. Given the permutation σ, let $\tilde{\sigma} = \rho_{A(\sigma)}(\sigma)$, and write $\tilde{\sigma}_i = \rho_{A_i(\sigma_i)}(\sigma_i)$ for the result of applying ρ_{A_i} to the ith cycle in the canonical decomposition of σ. Then

$$\begin{aligned}\rho_{A(\tilde{\sigma})}(\tilde{\sigma}) &= \prod_{i=1}^{K} \rho_{A_i(\tilde{\sigma}_i)}(\tilde{\sigma}_i) = \prod_{i=1}^{K} \rho_{A_i(\sigma_i)}(\tilde{\sigma}_i) \qquad \text{as in the proof of Lemma 4.7} \\ &= \prod_{i=1}^{K} \rho_{A_i(\sigma_i)}(\rho_{A_i(\sigma_i)}(\sigma_i)) = \prod_{i=1}^{K} \sigma_i = \sigma.\end{aligned}$$

So the function $\rho_{A(\cdot)}(\cdot) : \mathcal{S}_e \to \mathcal{S}_e$ is once again a bijection. □

We have described how to generate a new permutation $\tilde{\sigma}$ from a permutation σ such that it has the same law as σ. We now detail our method for choosing the vector $A(\sigma)$ appearing in the definition of ρ_A, in such a way that the conditions of Corollary 4.9 are satisfied. A will have the important property that all particles that are to be pinkened at a pinkening step in the chameleon process will be located at vertices in A at that time. Our choice of A will depend not only on σ but also on particular subsets of vertices in the edge under consideration. (Later on these subsets will be specified in the chameleon process, but for now we keep them as general subsets.) Indeed the function A is of the form $A(R, W, \sigma)$, where $(R, W, \sigma) \in \mathcal{P}^V \times \mathcal{P}^V \times \mathcal{S}_e$ for some $e \in E$.

Let R and W be two disjoint subsets of V. For each $i \in \{1, \dots, K\}$ (where recall that K is the number of cycles in σ) we define 8 sets:

$$\begin{aligned}
B_i^1(R, W, \sigma_i) &= \left\{ j \in [d_i'] : \sigma_i^{2j-2}(m_i) \in W,\, \sigma_i^{2d_i'+2j-2}(m_i) \in R \right\}, \\
B_i^2(R, W, \sigma_i) &= \left\{ j \in [d_i'] : \sigma_i^{2j-2}(m_i) \in R,\, \sigma_i^{2d_i'+2j-2}(m_i) \in W \right\}, \\
B_i^3(R, W, \sigma_i) &= \left\{ j \in [d_i'] : \sigma_i^{2j-2}(m_i) \in R,\, \sigma_i^{2d_i'+2j-2}(m_i) \in R \right\}, \\
B_i^4(R, W, \sigma_i) &= \left\{ j \in [d_i'] : \sigma_i^{2j-2}(m_i) \in W,\, \sigma_i^{2d_i'+2j-2}(m_i) \in W \right\}, \\
D_i^1(R, W, \sigma_i) &= \left\{ j \in [d_i'] : \sigma_i^{2j-1}(m_i) \in W,\, \sigma_i^{2d_i'+2j-1}(m_i) \in R \right\}, \\
D_i^2(R, W, \sigma_i) &= \left\{ j \in [d_i'] : \sigma_i^{2j-1}(m_i) \in R,\, \sigma_i^{2d_i'+2j-1}(m_i) \in W \right\}, \\
D_i^3(R, W, \sigma_i) &= \left\{ j \in [d_i'] : \sigma_i^{2j-1}(m_i) \in W,\, \sigma_i^{2d_i'+2j-1}(m_i) \in W \right\}, \\
D_i^4(R, W, \sigma_i) &= \left\{ j \in [d_i'] : \sigma_i^{2j-1}(m_i) \in R,\, \sigma_i^{2d_i'+2j-1}(m_i) \in R \right\}.
\end{aligned}$$

Note that (for each i) the B sets are disjoint, as are the D sets. Using these we define two further disjoint sets for each $i \in \{1, \dots, K\}$:

$$\begin{aligned}
C_i^1(R, W, \sigma_i) &= (B_i^1(R, W, \sigma_i) \cup B_i^2(R, W, \sigma_i)) \cap (D_i^3(R, W, \sigma_i) \cup D_i^4(R, W, \sigma_i)) \\
C_i^2(R, W, \sigma_i) &= (B_i^3(R, W, \sigma_i) \cup B_i^4(R, W, \sigma_i)) \cap (D_i^1(R, W, \sigma_i) \cup D_i^2(R, W, \sigma_i)).
\end{aligned}$$

For each $i \in \{1, \dots, K\}$, we define

$$A_i(R, W, \sigma_i) = C_i^1(R, W, \sigma_i) \cup C_i^2(R, W, \sigma_i)\,, \tag{4.7}$$

and then finally, we define $A(R, W, \sigma)$ to be the vector

$$A(R, W, \sigma) = (A_1(R, W, \sigma_1), A_2(R, W, \sigma_2), \dots, A_K(R, W, \sigma_K))\ . \tag{4.8}$$

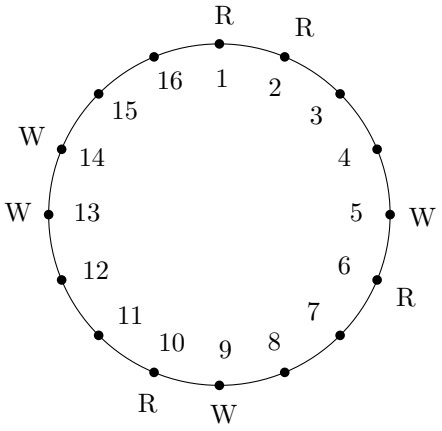

Figure 6: In this example we suppose $e = \{1, \ldots, 16\}$, σ is the cycle $(1\,2\,\cdots\,16)$, $R = \{1, 2, 6, 10\}$, and $W = \{5, 9, 13, 14\}$. Then the non-empty sets are $B^2 = D^4 = C^1 = \{1\}$, $B^4 = D^2 = C^2 = \{3\}$ and so $A(R, W, \sigma) = \{1, 3\}$.

We conclude this section by showing that this construction of A_i satisfies the property $A_i(R, W, \rho_{A_i(\sigma_i)}(\sigma_i)) = A_i(R, W, \sigma_i)$. This in turn will imply that the vector A satisfies the conditions of Corollary 4.9, which will guarantee that the permutation $\rho_A(\sigma)$ has the same law on $\mathcal{S}_e$ as σ.

Lemma 4.10. *Fix an edge $e \in E$. For any disjoint subsets R, W of V and permutation $\sigma \in \mathcal{S}_e$, the function $A_i = A_i(R, W, \sigma_i)$ defined in* (4.7) *satisfies*

$$A_i(R, W, \rho_{A_i(R,W,\sigma_i)}(\sigma_i)) = A_i(R, W, \sigma_i).$$

Proof. Suppose $k \in A_i(R, W, \sigma_i)$. It follows that either $k \in C_i^1(R, W, \sigma_i)$ or $k \in C_i^2(R, W, \sigma_i)$. We consider the first case, and further consider the subcase $k \in B_i^1(R, W, \sigma_i) \cap D_i^3(R, W, \sigma_i)$. By the definition of these two sets we see that

$$\begin{aligned} \sigma_i^{2k-2}(m_i) \in W, \quad \sigma_i^{2d_i'+2k-2}(m_i) \in R, \\ \sigma_i^{2k-1}(m_i) \in W, \quad \sigma_i^{2d_i'+2k-1}(m_i) \in W. \end{aligned}$$

Since $k \in A_i(R, W, \sigma_i)$, by the construction of $\rho_{A_i(\sigma_i)}(\sigma_i)$ in Definition 4.8 we deduce that

$$\begin{aligned} \rho_{A_i(\sigma_i)}(\sigma_i)^{2d_i'+2k-2}(m_i) = \sigma_i^{2k-2}(m_i) \in W, \\ \rho_{A_i(\sigma_i)}(\sigma_i)^{2k-2}(m_i) = \sigma_i^{2d_i'+2k-2}(m_i) \in R, \\ \rho_{A_i(\sigma_i)}(\sigma_i)^{2d_i'+2k-1}(m_i) = \sigma_i^{2k-1}(m_i) \in W, \\ \rho_{A_i(\sigma_i)}(\sigma_i)^{2k-1}(m_i) = \sigma_i^{2d_i'+2k-1}(m_i) \in W. \end{aligned}$$

Therefore

$$k \in B_i^2(R, W, \rho_{A_i(\sigma_i)}(\sigma_i)) \cap D_i^3(R, W, \rho_{A_i(\sigma_i)}(\sigma_i))$$

and so

$$k \in A_i(R, W, \rho_{A_i(\sigma_i)}(\sigma_i))$$

The other cases follow similarly. This shows that

$$A_i(R, W, \sigma_i) \subseteq A_i(R, W, \rho_{A_i(\sigma_i)}(\sigma_i)),$$

but an identical argument shows the reverse implication and we deduce the result. □

4.3 Chameleon process construction

We are now able to complete our formal definition of the chameleon process corresponding to an interchange process on a hypergraph. Let T be a (suitably large) integer, which shall be called the *phase length*. As stated previously, the chameleon process is time-inhomogeneous, and behaves differently depending on which phase we are in. There will be just two different kinds of phase: those in which no colour-changing is permitted and particles are just moved around the graph according to the underlying interchange process; and those in which colour-changing (pinkening of red and white particles) can occur. Furthermore, there will be (deterministic) times at which depinking can occur.

To be more precise, suppose $\mathcal{P}$ is the Poisson process used in the graphical construction of an interchange process, as in Section 2.2. Then the chameleon process is updated at the jump times $\{\tau_n\}$ of $\mathcal{P}$ and also at deterministic times $2iT$, $i \in \mathbb{N}$.

To describe which particles are pinkened during a colour-changing phase, let σ $(= \sigma_n)$ be the permutation applied to some edge e $(= e_n)$ at time $t = \tau_n$.

Given that t is in a colour-changing phase, we define subsets of V:

$$\begin{aligned} L_t^{i,1} &= \bigcup_{j \in C_i^1(R_{t-}, W_{t-}, \sigma_i)} \left\{\sigma_i^{2j-2}(m_i), \sigma_i^{2d_i'+2j-2}(m_i)\right\}, \\ L_t^{i,2} &= \bigcup_{j \in C_i^2(R_{t-}, W_{t-}, \sigma_i)} \left\{\sigma_i^{2j-1}(m_i), \sigma_i^{2d_i'+2j-1}(m_i)\right\}, \\ \text{and} \quad L_t &= \bigcup_{i=1}^{K} \left\{L_t^{i,1}, L_t^{i,2}\right\}, \end{aligned}$$

where recall that the sets C_i^1 and C_i^2 are used to define $A_i(R, W, \sigma_i)$ in (4.7). (In the above we have suppressed the dependence on n of quantities K, m_i, d_i etc, to ease notation.) Note that $|L_t \cap R_{t-}| = |L_t \cap W_{t-}| = |L_t/2|$. The particles at vertices in L_t are those that we wish to pinken at time t.

However, when changing the colours of particles we will not want the resulting number of pinks to exceed either the number of red or white particles. (This will make it easier in Section 5 to study how the number of pink particles evolves when viewed at depinking times.) In order to achieve this, for each $1 \le i \le K$, we (arbitrarily) choose subsets

$$\bar{C}_i^1(R_{t-}, W_{t-}, \sigma_i) \subseteq C_i^1(R_{t-}, W_{t-}, \sigma_i) \quad \text{and} \quad \bar{C}_i^2(R_{t-}, W_{t-}, \sigma_i) \subseteq C_i^2(R_{t-}, W_{t-}, \sigma_i)$$

satisfying

$$3\sum_{i=1}^{K} \left(|\bar{C}_i^1(R_{t-}, W_{t-}, \sigma_i)| + |\bar{C}_i^2(R_{t-}, W_{t-}, \sigma_i)| \right) \le \min\{|R_{t-}|, |W_{t-}|\} - |P_{t-}|,$$

and moreover we choose these subsets to be as large as possible so that the left-hand side in the above equation is maximal. We then define $\bar{L}_t^{i,1}$ and $\bar{L}_t^{i,2}$ to be as $L_t^{i,1}$ and $L_t^{i,2}$ replacing C_i^1 and C_i^2 with $\bar{C}_i^1$ and $\bar{C}_i^2$ respectively, and set

$$\bar{L}_t = \bigcup_{i=1}^{K} \{\bar{L}_t^{i,1}, \bar{L}_t^{i,2}\}.$$

It is precisely the particles at vertices in $\bar{L}_t$ that we will pinken at time t.

Box 4.11. Formal description of chameleon process updates
There are three kinds of updates to the chameleon process – which update is performed at time t depends on the value of t:

Constant-colour phases: For $t = \tau_n \in (2(i-1)T, (2i-1)T]$, update as the interchange process:

$$(\mathbf{z}_t, R_t, P_t, W_t) = (\sigma_n(\mathbf{z}_{t-}), \sigma_n(R_{t-}), \sigma_n(P_{t-}), \sigma_n(W_{t-})),$$

where for $U \subset V$,

$$\sigma_n(U) = \{\sigma_n(a) : a \in U\},$$

and for $\mathbf{v} \in (V)_{k-1}$,

$$\sigma_n(\mathbf{v}) = \big(\sigma_n(\mathbf{v}(1)), \ldots, \sigma_n(\mathbf{v}(k-1))\big).$$

Colour-changing phases: For $t = \tau_n \in ((2i-1)T, 2iT]$, update as interchange (i.e. update as in a constant-colour phase) unless $|P_{t-}| < \min\{|R_{t-}|, |W_{t-}|\}$. In this case, we pinken all particles in the set $\bar{L}_t$ (half of which belong to R_{t-}, the others to W_{t-}, by design).

Regardless of which particles are pinkened, we then update as the interchange process at this time (using σ_n). Formally then, we update as

$$(\mathbf{z}_t, R_t, P_t, W_t) = (\sigma_n(\mathbf{z}_{t-}), \sigma_n(R_{t-} \setminus (\bar{L}_t \cap R_{t-})), \sigma_n(P_{t-} \cup \bar{L}_t), \sigma_n(W_{t-} \setminus (\bar{L}_t \cap W_{t-}))).$$

Depinking: For $t = 2iT$ with $i \in \mathbb{N}$, if $|P_{t-}| \geq \min\{|R_{t-}|, |W_{t-}|\}$ then we toss a fair coin c_i: if c_i is heads, we colour all pink particles red, otherwise we colour all pink particles white. Hence the update is

$$(\mathbf{z}_t, R_t, P_t, W_t) = \begin{cases} (\mathbf{z}_{t-}, R_{t-} \cup P_{t-}, \varnothing, W_{t-}) & \text{if } c_i = 1, \\ (\mathbf{z}_{t-}, R_{t-}, \varnothing, W_{t-} \cup P_{t-}) & \text{if } c_i = 0. \end{cases}$$

In Section 4.1 we outlined the desired properties of the chameleon process, and showed that if these hold we can relate the location of the kth particle in an interchange process to the amount of ink in the chameleon process (Lemma 4.2). We conclude this section by verifying that the chameleon process that we have just constructed does indeed satisfy these desired properties. We recall the notation $\{\tau_n\}$ and $\{\bar{\tau}_n\}$ from Section 4.1.

Lemma 4.12 (Desired properties hold)**.**
Fix $s > 0$, $\mathbf{z} \in (V)_{k-1}$ and $M_0 = (\mathbf{z}, R, P, W) \in \Omega_k(V)$. Let $\{M_t\}_{t \geq s}$ be a chameleon process as described in Box 4.11 with $M_s := M_0$, defined with respect to some interchange process $\{\mathbf{x}_t^{\mathrm{IP}}\}_{t \geq s}$ with $\mathbf{x}_s^{\mathrm{IP}} := \mathbf{x} = (\mathbf{z}, x) \in (V)_k$. Then all of the desired properties in Box 4.1 are satisfied.

Proof. We show that the desired properties hold at time $\bar{\tau}_1^s$, whether this corresponds to an update of the interchange process or is a depinking time. (Numbering within the proof corresponds to that of the properties in Box 4.1.)

1. Suppose that $\bar{\tau}_1^s = \tau_1^s$ (i.e. the first update of M corresponds to an incident in $\mathbf{x}^{\mathrm{IP}}$). In the case that τ_1^s occurs in a constant-colour phase, we simply set $\tilde{\sigma}_1^s = \sigma_1^s$, so that $\{\mathbf{x}_t^{\mathrm{IP}}\} = \{\tilde{\mathbf{x}}_t^{\mathrm{IP}}\}$ (and so properties 1a and 1b are satisfied). Property 1c follows since for every $v \in V$,
$$\mathrm{ink}_{s+\tau_1^s}(v) = \mathrm{ink}_s((\sigma_1^s)^{-1}(v)).$$

   Now suppose τ_1^s is in a colour-changing phase. We set $\tilde{\sigma}_1^s = \rho_{A(R,W,\sigma_1^s)}(\sigma_1^s)$. Property 1a follows from the fact that ρ_A only affects the movement of red and white particles (by Part 3 of Lemma 4.6 and the construction of A in Section 4.3). Furthermore, property 1b is satisfied thanks to Lemma 4.10 and Corollary 4.9, which show that $\tilde{\sigma}_1^s$ has the same law as σ_1^s and is independent of R and W.

   It remains to prove property 1c when τ_1^s is in a colour-changing phase. For ease of notation we write σ for σ_1^s and $\tilde{\sigma}$ for $\tilde{\sigma}_1^s$, and use the notation σ_i, m_i, d_i etc from Section 4.2. Suppose that v satisfies
$$\mathrm{ink}_s(\sigma^{-1}(v)) \neq \mathrm{ink}_s(\tilde{\sigma}^{-1}(v)).$$

   By the construction of A in (4.8), this is true if and only if there exists $i \in \{1, \ldots, K(\sigma)\}$ such that one of the following holds:

   - $\exists \ell \in \bar{C}_i^1(R, W, \sigma_i)$ with $v \in \{\sigma_i^{2\ell-2}(m_i), \sigma_i^{2d_i'+2\ell-2}(m_i)\}$,
   - $\exists \ell \in \bar{C}_i^2(R, W, \sigma_i)$ with $v \in \{\sigma_i^{2\ell-1}(m_i), \sigma_i^{2d_i'+2\ell-1}(m_i)\}$.

   However, this is equivalent to the particle at v being pinkened at time $s + \tau_1^s$ in the chameleon process M. Thus almost surely $\mathrm{ink}_{s+\tau_1^s}(v) = 1/2$, and almost surely,
$$\{\mathrm{ink}_s(\sigma^{-1}(v)), \mathrm{ink}_s(\tilde{\sigma}^{-1}(v))\} = \{0, 1\}.$$

2. If $\bar{\tau}_1^s$ is a depinking time, then $\mathbb{E}[\mathrm{ink}_{\bar{\tau}_1^s}(v)] = \mathrm{ink}_0(v)$ for all $v \in V$ thanks to the fair coin toss used to decide whether pink particles are coloured red or white. Thus property 2 is satisfied.

□

5 From k particles to 4: proof of Lemma 1.8

We complete the proof of Lemma 1.8 in this section by using the chameleon process to bound the mixing time of the interchange process. In order to do this, we will need to understand the evolution of the total amount of ink in the chameleon process. We first of all note that the number of pink particles accumulates over time until we have a large number of them, at which point all pink particles are recoloured (either red or white) and the process of accumulation starts again. The process will continue in this manner until either we have no white particles or we have no red particles (which will occur immediately after some depinking). At this point, no more pink particles can be made and so there is no more recolouring of particles. In order to bound the mixing time of the interchange process we need a good understanding of how quickly the chameleon process reaches the state where no more recolouring can occur. There are two factors which affect this: the time we must wait between depinking events and how the process behaves at depinking times.

Writing $\mathbf{x} = (\mathbf{z}, x)$, for each $j \in \mathbb{N}$ let $D_j(\mathbf{x})$ denote the jth depinking time of a chameleon process started from state $(\mathbf{z}, \{x\}, \varnothing, V \setminus (\mathbf{O}(\mathbf{z}) \cup \{x\})) \in \Omega_k(V)$. Let $\mathrm{ink}_t^{\mathbf{x}}$ denote the total amount of ink in the process at time t; note that $0 \leq \mathrm{ink}_t^{\mathbf{x}} \leq |V| - k + 1$. Motivated by Oliveira (2013), we define the event

$$\mathrm{Fill}^{\mathbf{x}} := \left\{ \lim_{t\to\infty} \mathrm{ink}_t^{\mathbf{x}} = |V| - k + 1 \right\}.$$

This is the event that all initially-white particles are eventually coloured red. We shall make use of the following result concerning $\mathrm{ink}_t^{\mathbf{x}}$, whose proof may be found in Oliveira (2013).

Proposition 5.1. *Fix $\mathbf{x} = (\mathbf{z}, x) \in (V)_k$. For each $\mathbf{c} \in (V)_{k-1}$ and $t \geq 0$,*

$$\mathbb{P}[\{\mathbf{z}_t = \mathbf{c}\} \cap \mathrm{Fill}^{\mathbf{x}}] = \frac{\mathbb{P}[\mathbf{z}_t = \mathbf{c}]}{|V| - k + 1}.$$

We can now relate the total-variation distance between two realisations of $\mathrm{IP}(k, f, G)$ to a certain expectation involving the amount of ink in a chameleon process. The following result is similar to Lemma 6.1 of Oliveira (2013): we include a sketch of the proof to highlight the importance of constructing in Section 4 a chameleon process satisfying Lemma 4.2.

Lemma 5.2. *For every $t \geq 0$,*

$$\sup_{\mathbf{x},\mathbf{y}\in(V)_k} \|\mathcal{L}[\mathbf{x}_t^{\mathrm{IP}}] - \mathcal{L}[\mathbf{y}_t^{\mathrm{IP}}]\|_{\mathrm{TV}} \leq 2k \sup_{\mathbf{w}\in(V)_k} \mathbb{E}\left[1 - \frac{\mathrm{ink}_t^{\mathbf{w}}}{|V| - k + 1}\Big| \,\mathrm{Fill}^{\mathbf{w}}\right]$$

Proof. Fix $\mathbf{x} = (\mathbf{z}, x) \in (V)_k$ with $\mathbf{z} \in (V)_{k-1}$, and denote by $\mathbf{x}_t^{\mathrm{IP}}$ an interchange process started from $\mathbf{x}$. Let $\tilde{x}$ be uniform from $V \setminus \mathbf{O}(\mathbf{z})$ and denote by $\tilde{\mathbf{x}}_t^{\mathrm{IP}}$ an interchange process started from $\tilde{\mathbf{x}} = (\mathbf{z}, \tilde{x})$. Then for any $\mathbf{b} = (\mathbf{c}, b) \in (V)_k$,

$$\mathbb{P}[\tilde{\mathbf{x}}_t^{\mathrm{IP}} = \mathbf{b}] = \frac{\mathbb{P}[\mathbf{z}_t^{\mathrm{IP}} = \mathbf{c}]}{|V| - k + 1} = \frac{\mathbb{P}[\mathbf{z}_t = \mathbf{c}]}{|V| - k + 1} = \mathbb{P}[\{\mathbf{z}_t = \mathbf{c}\} \cap \mathrm{Fill}^{\mathbf{x}}],$$

where the second equality follows from the fact that the black particles in a chameleon process move according to the underlying interchange process (with no recolouring), and the third from Proposition 5.1. On the other hand, thanks to Lemma 4.2:

$$\mathbb{P}[\mathbf{x}_t^{\mathrm{IP}} = \mathbf{b}] = \mathbb{E}[\mathrm{ink}_t(b)\mathbf{1}_{\{\mathbf{z}_t=\mathbf{c}\}}] \geq \mathbb{E}[\mathrm{ink}_t(b)\mathbf{1}_{\{\{\mathbf{z}_t=\mathbf{c}\}\cap\mathrm{Fill}^{\mathbf{x}}\}}].$$

Subtracting we obtain

$$\mathbb{P}[\tilde{\mathbf{x}}_t^{\mathrm{IP}} = \mathbf{b}] - \mathbb{P}[\mathbf{x}_t^{\mathrm{IP}} = \mathbf{b}] \leq \mathbb{E}[(1 - \mathrm{ink}_t(b))\mathbf{1}_{\{\{\mathbf{z}_t = \mathbf{c}\} \cap \mathrm{Fill}^{\mathbf{x}}\}}].$$

Hence

$$\begin{aligned}\|\mathcal{L}[\mathbf{x}_t^{\mathrm{IP}}] - \mathcal{L}[\tilde{\mathbf{x}}_t^{\mathrm{IP}}]\|_{\mathrm{TV}} &\leq \sum_{(\mathbf{c},b) \in (V)_k} \mathbb{E}[(1 - \mathrm{ink}_t(b))\mathbf{1}_{\{\{\mathbf{z}_t = \mathbf{c}\} \cap \mathrm{Fill}^{\mathbf{x}}\}}] \\ &= \mathbb{E}[(|V| - k + 1 - \mathrm{ink}_t^{\mathbf{x}})\mathbf{1}_{\{\mathrm{Fill}^{\mathbf{x}}\}}] \\ &= \mathbb{E}\left[1 - \frac{\mathrm{ink}_t^{\mathbf{x}}}{|V| - k + 1} \middle| \, \mathrm{Fill}^{\mathbf{x}}\right].\end{aligned}$$

The result now follows by repeated application of the triangle inequality, as in the proof of Lemma 6.1 of Oliveira (2013). □

Consider now the expectation on the right-hand side of the statement of Lemma 5.2: an identical argument to that in Section 6 of Oliveira (2013) shows that this can be bounded in terms of the tail probability of the time of the jth depinking.

Lemma 5.3. *There exist positive constants c_1 and c_2 such that for every $j \in \mathbb{N}$,*

$$\sup_{\mathbf{x} \in (V)_k} \mathbb{E}\left[1 - \frac{\mathrm{ink}_t^{\mathbf{x}}}{|V| - k + 1} \middle| \, \mathrm{Fill}^{\mathbf{x}}\right] \leq c_1 \sqrt{|V|} e^{-c_2 j} + \sup_{\mathbf{x} \in (V)_k} \mathbb{P}[D_j(\mathbf{x}) > t | \, \mathrm{Fill}^{\mathbf{x}}]$$

We therefore see that we need good control on the probability that there have only been a few depinkings by time t. Here we cannot simply rely on results from Oliveira (2013), since the chameleon process constructed in Section 4 clearly obeys different dynamics. We shall need the following fundamental result – a lower bound on the number of red particles that are lost (due to pinkening) during a colour-changing phase of the chameleon process (where we start the phase with more white particles than red). The proof is deferred to Section 6.

Lemma 5.4. *Consider a chameleon process with initial configuration $(\mathbf{z}, R, P, W)$ satisfying $|P| < |R| \leq |W|$. If the phase length parameter T satisfies $T \geq 20 T_{\mathrm{IP}(4,f,G)}(1/4)$, then*

$$\mathbb{E}[|R_{2T-}|] \leq (1 - 10^{-6})|R|.$$

This result allows us to bound the probability appearing in the statement of Lemma 5.3.

Lemma 5.5. *There exists a universal constant $\kappa_1 > 0$ such that for every interchange process with G and f satisfying Assumptions 1.1 and 1.2 respectively, every $j \in \mathbb{N}$ and $\mathbf{x} \in (V)_k$,*

$$\mathbb{P}[D_j(\mathbf{x}) > t | \, \mathrm{Fill}^{\mathbf{x}}] \leq \exp\left\{ j - \frac{t}{\kappa_1 \, T_{\mathrm{IP}(4,f,G)}(1/4)} \right\}.$$

Proof. Thanks to Lemma 5.4, the proofs of Lemmas 6.2 and 9.2 of Oliveira (2013) can be emulated to show that there exists a positive constant κ such that

$$\mathbb{E}[e^{D_j(\mathbf{x})/\kappa T} | \, \mathrm{Fill}^{\mathbf{x}}] \leq e^j,$$

for all $j \in \mathbb{N}$, where $T = 20 T_{\mathrm{IP}(4,f,G)}(1/4)$. Thus by Markov's inequality,

$$\mathbb{P}[D_j(\mathbf{x}) > t | \, \mathrm{Fill}^{\mathbf{x}}] = \mathbb{P}[e^{D_j(\mathbf{x})/\kappa T} > e^{t/\kappa T} | \, \mathrm{Fill}^{\mathbf{x}}] \leq e^{-t/\kappa T} \mathbb{E}[e^{D_j(\mathbf{x})/\kappa T} | \, \mathrm{Fill}^{\mathbf{x}}] \leq e^{j - t/\kappa T}.$$

Writing $\kappa_1 = 20\kappa$ completes the proof. □

We are now finally able to complete the proof of Lemma 1.8, and thus bound the mixing time of $\mathrm{IP}(k, f, G)$ by that of $\mathrm{IP}(4, f, G)$.

Proof of Lemma 1.8. We start by combining Lemmas 5.2, 5.3 and 5.5 to give for every $t \geq 0$ and $j \in \mathbb{N}$,

$$\sup_{\mathbf{x},\mathbf{y}\in(V)_k} \|\mathcal{L}[\mathbf{x}_t^{\mathrm{IP}}] - \mathcal{L}[\mathbf{y}_t^{\mathrm{IP}}]\|_{\mathrm{TV}} \leq 2k \left\{ c_1\sqrt{|V|}e^{-c_2 j} + \exp\left\{ j - \frac{t}{\kappa_1 T_{\mathrm{IP}(4,f,G)}(1/4)} \right\} \right\},$$

for some universal positive constants c_1, c_2 and κ_1. We choose

$$j = \left\lfloor \frac{t}{(1+c_2)\kappa_1 T_{\mathrm{IP}(4,f,G)}(1/4)} \right\rfloor,$$

which gives the bound (using $k \leq |V|$),

$$\sup_{\mathbf{x},\mathbf{y}\in(V)_k} \|\mathcal{L}[\mathbf{x}_t^{\mathrm{IP}}] - \mathcal{L}[\mathbf{y}_t^{\mathrm{IP}}]\|_{\mathrm{TV}} \leq c_3|V|^{3/2} \exp\left\{ \frac{-c_2 t}{(1+c_2)\kappa_1 T_{\mathrm{IP}(4,f,G)}(1/4)} \right\},$$

for some positive c_3. Therefore there exists a universal constant C such that for any $\varepsilon \in (0, 1/2)$ and $t > C T_{\mathrm{IP}(4,f,G)}(1/4)\log(|V|/\varepsilon)$,

$$\sup_{\mathbf{x},\mathbf{y}\in(V)_k} \|\mathcal{L}[\mathbf{x}_t^{\mathrm{IP}}] - \mathcal{L}[\mathbf{y}_t^{\mathrm{IP}}]\|_{\mathrm{TV}} \leq \varepsilon. \qquad \square$$

6 Counting the number of pinkenings: proof of Lemma 5.4

In order to prove Lemma 5.4, we need to show that during a colour-changing phase (started with more white particles than red particles) the number of pink particles we create is in expectation at least a constant times the number of red particles at the start of that phase. We prove this result in this section.

Suppose we wish to lower-bound the number of pinkenings during the first colour-changing phase $[T, 2T]$. Since we start with 1 red particle, there will be more white particles at time T than red. We will wish to apply the following analysis for a general colour-changing phase (and not just the first) but the calculations carry through provided the number of white particles is at least the number of red. If this is not the case at the start of some colour-changing phase, then we can simply reverse the roles of the white and red particles in the following.

Recall that we initialise our chameleon process from configuration $(\mathbf{z}, R, P, W)$ and so W denotes the set of vertices occupied by white particles. We now choose W^r to be an arbitrary subset of W of size $\lfloor |W|/3 \rfloor$ and think of W^r as a *reserved* set of white particles, in the sense that we will not allow them to be pinkened during the evolution of the chameleon process. Of course, this defines a new process but it is clear that making this change cannot increase the number of pinkening events, and it turns out that this modification makes counting the expected number of pinkening events easier.

We make one further change to the chameleon process in this section in order to ease our analysis – we remove the condition that we only pinken if we have fewer pink particles than

either red or white particles. Although this means we can end up with more pinkening events, this will only happen if a certain number of pinkening events have already happened (we can only end up with having many pink particles through having lots of pinkening events), and in that case we will be happy regardless.

To formalise these two modifications we define

$$\tilde{B}_i^4(R,W,\sigma_i) = \big\{j \in [d_i'] : \sigma_i^{2j-2}(m_i) \in W^r,\, \sigma_i^{2d_i'+2j-2}(m_i) \in W^r\big\},$$
$$\tilde{D}_i^3(R,W,\sigma_i) = \big\{j \in [d_i'] : \sigma_i^{2j-1}(m_i) \in W^r,\, \sigma_i^{2d_i'+2j-1}(m_i) \in W^r\big\},$$

and then set

$$\tilde{C}_i^1(R,W,\sigma_i) = (B_i^1(R,W,\sigma_i) \cup B_i^2(R,W,\sigma_i)) \cap (\tilde{D}_i^3(R,W,\sigma_i),$$
$$\tilde{C}_i^2(R,W,\sigma_i) = \tilde{B}_i^4(R,W,\sigma_i) \cap (D_i^1(R,W,\sigma_i) \cup D_i^2(R,W,\sigma_i)).$$

We modify the colour-changing phase by exchanging the C_i^1 and C_i^2 appearing in Section 4.3 with $\tilde{C}_i^1$ and $\tilde{C}_i^2$ and removing the condition $|P_{t-}| < \min\{|R_{t-}|, |W_{t-}|\}$. We shall refer to this new process as the *modified chameleon process*.

Let $a \in V$. We define t_a to be the smallest integer n such that $T < \tau_n \leq 2T$ and $a \in e_n$. If no such n exists we set $t_a = \infty$. Also, we set $\phi_a = \tau_{t_a}$ with notation $\tau_\infty = \infty$; hence ϕ_a is the first time (after time T) that vertex a is in a ringing edge of the underlying interchange process. We define a third variable, F_a, set to be equal to $*$ in the case $\phi_a = \infty$. If, on the other hand, $\phi_a < \infty$, there are two possible cases. Let c_{t_a} denote the cycle of σ_{t_a} containing vertex a.

1. If $|c_{t_a}| \geq 4$, and there exists $j \in \{1, \ldots, d'(c_{t_a})\}$ with

   (i). $a = c_{t_a}^{2j-2}(m(c_{t_a}))$, then set

   $$F_a = I_{[T,\phi_a)}^{-1}\big(c_{t_a}^{2d'(c_{t_a})+2j-2}(m(c_{t_a})), c_{t_a}^{2j-1}(m(c_{t_a})), c_{t_a}^{2d'(c_{t_a})+2j-1}(m(c_{t_a}))\big);$$

   (ii). $a = c_{t_a}^{2d'(c_{t_a})+2j-2}(m(c_{t_a}))$, then set

   $$F_a = I_{[T,\phi_a)}^{-1}\big(c_{t_a}^{2j-2}(m(c_{t_a})), c_{t_a}^{2j-1}(m(c_{t_a})), c_{t_a}^{2d'(c_{t_a})+2j-1}(m(c_{t_a}))\big);$$

   (iii). $a = c_{t_a}^{2j-1}(m(c_{t_a}))$, then set

   $$F_a = I_{[T,\phi_a)}^{-1}\big(c_{t_a}^{2d'(c_{t_a})+2j-1}(m(c_{t_a})), c_{t_a}^{2j-2}(m(c_{t_a})), c_{t_a}^{2d'(c_{t_a})+2j-2}(m(c_{t_a}))\big);$$

   (iv). $a = c_{t_a}^{2d'(c_{t_a})+2j-1}(m(c_{t_a}))$, then set

   $$F_a = I_{[T,\phi_a)}^{-1}\big(c_{t_a}^{2j-1}(m(c_{t_a})), c_{t_a}^{2j-2}(m(c_{t_a})), c_{t_a}^{2d'(c_{t_a})+2j-2}(m(c_{t_a}))\big).$$

2. if $|c_{t_a}| \geq 4$ but there is no such j as above, or $|c_{t_a}| < 4$, then set $F_a = *$.

We denote the vertices containing non-reserved white particles at time 0 by $W^{\complement}$, i.e. $W^{\complement} = W \setminus W^r$ (recall that it is only these particles that we allow the be pinkened in the modified chameleon process). We now present a method to count the number of pinkenings during a colour-changing phase of the modified chameleon process. For ease of notation we shall write I for the map $I_{[0,T]}$. The proofs of the first three results below can be found in Appendix B.

Lemma 6.1. *Consider a modified chameleon process with starting configuration $(\mathbf{z}, R, P, W)$. Then the number of pinkenings during $(T, 2T)$ is at least the number of $b \in I(W^{\complement})$ such that $F_a = (b, c, d)$ for some $a \in I(R)$ and $c, d \in I(W^r)$ with $\phi_a = \phi_b$.*

In bounding the expected number of pinkenings during a colour-changing phase, it turns out to be useful to have a lower bound on the probability that $F_a \neq *$ given $\phi_a \neq \infty$. This is because even if vertex a is in a ringing edge during time interval $[T, 2T]$, in order for the particle initially at a to be pinkened in the modified chameleon process at this time, it is necessary (but not sufficient) that $F_a \neq *$.

Lemma 6.2. *For every $a \in V$,*

$$\mathbb{P}[F_a \neq *|\, \phi_a \neq \infty] \geq 3/7.$$

Proposition 6.3. *Consider a modified chameleon process with initial configuration $(\mathbf{z}, R, P, W)$. If the phase length parameter T satisfies $T \geq 20T_{\mathrm{IP}(4,f,G)}(1/4)$, then for any vertices a, b, c, d,*

$$\mathbb{P}[a \in I(R), b \in I(W^{\complement}), c, d \in I(W^r)] \geq (1 - 2^{-9})^2 \frac{|R||W^{\complement}||W^r|^2}{|V|^4}.$$

We now present the main result of this section – a version of Lemma 5.4 but proved for the *modified* chameleon process. As explained earlier in this section, this implies the corresponding result for our original chameleon process.

Lemma 6.4. *Consider a modified chameleon process with initial configuration $(\mathbf{z}, R, P, W)$ satisfying $|P| < |R| \leq |W|$. If the phase length parameter T satisfies $T \geq 20T_{\mathrm{IP}(4,f,G)}(1/4)$, then*

$$\mathbb{E}[|R_{2T-}|] \leq (1 - 10^{-6})|R|.$$

Proof. Write $N(b)$ for the set of vertices that share at least one edge with b. By Lemma 6.1, we have

$$\begin{aligned} |R_{2T-}| &\leq |R| - \sum_{b \in I(W^{\complement})} \mathbf{1}\left\{ \bigcup_{\substack{a \in N(b) \\ c,d \in I(W^r)}} \{F_a = (b,c,d), \phi_a = \phi_b,\, a \in I(R)\} \right\} \\ &= |R| - \sum_{b \in I(W^{\complement})} \sum_{\substack{a \in N(b) \\ c,d \in I(W^r)}} \mathbf{1}_{\{F_a = (b,c,d), \phi_a = \phi_b,\, a \in I(R)\}}. \end{aligned}$$

Taking expectations the right-hand side becomes

$$\sum_b \sum_{\substack{a \in N(b) \\ c,d}} \mathbb{P}[F_a = (b,c,d), \phi_a = \phi_b,\, a \in I(R), b \in I(W^{\complement}), c, d \in I(W^r)].$$

Note now that the event $\{F_a = (b,c,d), \phi_a = \phi_b\}$ is determined entirely by the process after time T, and in particular is independent of the process between times 0 and T, and hence of the map $I = I_{[0,T]}$. Using Proposition 6.3 the above becomes

$$\begin{aligned}
&\sum_b \sum_{\substack{a \in N(b) \\ c,d}} \mathbb{P}[F_a = (b,c,d), \phi_a = \phi_b] \mathbb{P}[a \in I(R), b \in I(W^{\complement}), c,d \in I(W^r)] \\
&\geq \sum_b \sum_{\substack{a \in N(b) \\ c,d}} \mathbb{P}[F_a = (b,c,d), \phi_a = \phi_b] (1-2^{-9})^2 \frac{|R||W^{\complement}||W^r|^2}{|V|^4} \\
&= (1-2^{-9})^2 \frac{|R||W^{\complement}||W^r|^2}{|V|^4} \sum_b \sum_{a \in N(b)} \mathbb{P}[F_a(1) = b, \phi_a = \phi_b] \qquad (6.1)
\end{aligned}$$

where we write $F_a(1)$ for the first coordinate of F_a. Now

$$\mathbb{P}[F_a(1) = b, \phi_a = \phi_b] = \mathbb{P}[F_a(1) = b, \phi_a = \phi_b |\, F_a \neq *]\, \mathbb{P}[F_a \neq *]. \qquad (6.2)$$

Using Lemma 6.2,

$$\mathbb{P}[F_a \neq *] = \mathbb{P}[\phi_a \neq \infty] \mathbb{P}[F_a \neq * |\, \phi_a \neq \infty] \geq \frac{3}{7} \mathbb{P}[\phi_a \neq \infty]. \qquad (6.3)$$

Also,

$$\mathbb{P}[\phi_a = \infty] \leq \mathbb{P}[I_{(T,2T)}(a) = a] = \mathbb{P}[a_T^{\mathrm{RW}} = a],$$

where $\{a_t^{\mathrm{RW}}\}$ is a realisation of $\mathrm{RW}(f,G)$ started from a. Since

$$T = 20 T_{\mathrm{IP}(4,f,G)}(1/4) \geq 20 T_{\mathrm{RW}(f,G)}(1/4) \geq T_{\mathrm{RW}(f,G)}(2^{-20})$$

(by Proposition 2.2) we have

$$\mathbb{P}[\phi_a = \infty] \leq \mathbb{P}[a_T^{\mathrm{RW}} = a] \leq \frac{1}{|V|} + 2^{-20}.$$

From (6.3), we deduce

$$\mathbb{P}[F_a \neq *] \geq \frac{3}{7}(1 - 2^{-20} - 1/|V|). \qquad (6.4)$$

We now lower bound $\mathbb{P}[F_a(1) = b, \phi_a = \phi_b |\, F_a \neq *]$. Since all edges in the interchange process ring at equal rate, and f is constant on conjugacy classes (Assumption 1.2)

$$\begin{aligned}
\mathbb{P}[F_a(1) = b, \phi_a = \phi_b |\, F_a \neq *] &= \sum_{e : a,b \in e} \mathbb{P}[F_a(1) = b, \phi_a = \phi_b, e_{t_a} = e |\, F_a \neq *] \\
&= \sum_{e : a,b \in e} \big((|e|-1)(\deg(a) + \deg(b) - \deg(a,b)\big)^{-1}, \\
&\geq \sum_{e : a,b \in e} \big((|e|-1)(\deg(a) + \deg(b)\big)^{-1}. \qquad (6.5)
\end{aligned}$$

where $\deg(a,b)$ denotes the number of edges that a and b have in common. Now, by Assumption 1.1,

$$\sum_b \sum_{a \in N(b)} \sum_{e : a,b \in e} \frac{1}{(|e|-1)(\deg(a) + \deg(b))} = \sum_{e \in E} \sum_{a,b \in \binom{e}{2}} \frac{2}{(|e|-1)(\deg(a) + \deg(b))} \geq \frac{|V|}{10}.$$

Using this together with (6.1), (6.2), (6.4) and (6.5) we have

$$\mathbb{E}[|R_{2T-}| - |R|] \le -\frac{3}{7}\big(1 - 2^{-20} - 1/|V|\big)(1-2^{-9})^2 \frac{|R||W^{\mathtt{C}}||W^r|^2}{|V|^4}\frac{|V|}{10}.$$

To complete the proof, we recall that $|V| \ge 36$ (Assumption 1.1), that $|W^r| = \lfloor |W|/3 \rfloor$ and $|W^{\mathtt{C}}| = |W| - |W^r|$, and note that

$$3|W| \ge |W| + |R| + |P| = |V| - k + 1 \ge |V|/2\,.$$

Hence

$$\begin{aligned}\mathbb{E}[|R_{2T-}| - |R|] &\le -\frac{3}{7}\big(1 - 2^{-20} - 1/|V|\big)(1-2^{-9})^2 \frac{|R|}{10}(1/9)(1/18)^2 \\ &\le -10^{-6}|R|\,.\end{aligned}$$

□

A Technical proofs for Section 3

Here we include some of the more technical proofs required to compare the mixing time of $\mathrm{IP}(4, f, G)$ with that of $\mathrm{EX}(2, f, G)$.

Proof of Lemma 3.4. Let $\hat{\Lambda}$ be a Poisson process of rate $2|E|$ (i.e. double the usual rate), and let $\{e_n\}_{n\in\mathbb{N}}$ and $\{\sigma_n\}_{n\in\mathbb{N}}$ be random sequences of edges and permutations as described in Section 2.2. Let $\{c_n\}_{n\in\mathbb{N}}$ be a sequence of i.i.d. Bernoulli(1/2) random variables.

Call an edge e *special* at time $t \in [(i-1)\kappa T, i\kappa T)$ if it contains both $\mathbf{x}^{\mathrm{IP}}_{t-}(i)$ and $\mathbf{y}^{\mathrm{IP}}_{t-}(i)$, with $\mathbf{x}^{\mathrm{IP}}_{t-}(i) \ne \mathbf{y}^{\mathrm{IP}}_{t-}(i)$. Note that a special edge may contain some already matched particles. We couple $\{\mathbf{x}^{\mathrm{IP}}_t\}$ and $\{\mathbf{y}^{\mathrm{IP}}_t\}$ during time interval $[(i-1)\kappa T, i\kappa T)$ so that they evolve according to the following rules: at the time $\hat{\tau}_n \in [(i-1)\kappa T, i\kappa T)$ of the n^{th} incident of the Poisson process $\hat{\Lambda}$, if $c_n = 0$ then we do nothing, but if $c_n = 1$ then

1. if e_n is not special at time $\hat{\tau}_n$, we apply the permutation σ_n to both $\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}_n-}$ and $\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}_n-}$;
2. if e_n is special, we apply permutation σ_n to $\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}_n-}$, and construct a new permutation $\tilde{\sigma}_n$ to apply to $\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}_n-}$, as follows. Let U_n be the set of vertices in e_n excluding $\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}_n-}(i)$ and $\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}_n-}(i)$ and any vertices occupied by matched pairs of particles.

   (i). if $\sigma_n^{-1}(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}_n-}(i)) \in U_n$ *and* $\sigma_n^{-1}(\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}_n-}(i)) \in U_n$ then set

$$\tilde{\sigma}_n = (\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}_n-}(i)\, \mathbf{y}^{\mathrm{IP}}_{\hat{\tau}_n-}(i))(\sigma_n^{-1}(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}_n-}(i))\, \sigma_n^{-1}(\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}_n-}(i))) \circ \sigma_n\,.$$

   (Note that for permutations σ and ρ, we use the convention that their composition corresponds to multiplication on the right: $\rho \circ \sigma = \sigma\rho$.)
   This makes $\tilde{\sigma}_n(\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}_n-}(i)) = \sigma_n(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}_n-}(i))$, and so particle pair i become matched after these permutations are applied. Furthermore, any matched particles on our special edge are moved identically under $\tilde{\sigma}_n$ and σ_n, and so no matches are broken;

(ii). otherwise, set $\tilde{\sigma}_n = \sigma_n$.

We denote by $\hat{\tau}^*$ the first time after $(i-1)\kappa T$ that an edge is special, and write e^* for this edge and n^* for the integer satisfying $\hat{\tau}_{n^*} = \hat{\tau}^*$. We claim that for each n, $\tilde{\sigma}_n$ has the same distribution as σ_n (and so this really is a coupling of two interchange processes) and also

$$\mathbb{P}[\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*}(i) = \mathbf{y}^{\mathrm{IP}}_{\hat{\tau}^*}(i)] \geq \frac{(d-4)(d-5)(d-6)}{2(d-1)(d-2)(d-3)}\,,$$

where d is the minimum edge size. Indeed, the first claim is immediate from the construction of $\tilde{\sigma}_n$: by inspection $\tilde{\sigma}$ and σ belong to the same conjugacy class and moreover there is a bijection between σ_n and $\tilde{\sigma}_n$ since

$$\tilde{\sigma}_n^{-1}(\{\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}_n-}(i), \mathbf{y}^{\mathrm{IP}}_{\hat{\tau}_n-}(i)\}) = \sigma_n^{-1}(\{\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}_n-}(i)\}, \mathbf{x}^{\mathrm{IP}}_{\hat{\tau}_n-}(i)\}) \subseteq U_n\,.$$

For the second claim, we have,

$$\begin{aligned}\mathbb{P}[\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*}(i) = \mathbf{y}^{\mathrm{IP}}_{\hat{\tau}^*}(i)] &= \mathbb{P}[c_{n^*} = 1] \times \mathbb{P}[\sigma_{n^*}^{-1}(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \in U_{n^*},\, \sigma_{n^*}^{-1}(\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \in U_{n^*}] \\ &= \frac{1}{2}\,\mathbb{P}[\sigma_{n^*}^{-1}(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \in U_{n^*},\, \sigma_{n^*}^{-1}(\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \in U_{n^*}]\,. \end{aligned} \tag{A.1}$$

In order to bound the probability appearing on the right-hand side, note that since the special edge e^* contains $\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)$ and $\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)$, and there can be at most three pairs of particles already matched (regardless of the value of i) in e^*, we must have $|U_{n^*}| \geq d-5 \geq 2$, by Assumption 1.1. For the permutation σ chosen at time n^*, let $N_\sigma(1)$ be the number of fixed points in σ, and let us write $C_\sigma(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i))$ for the length of the cycle of σ containing $\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)$. Then:

$$\begin{aligned}&\mathbb{P}[\sigma_{n^*}^{-1}(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \in U_{n^*}\,,\, \sigma_{n^*}^{-1}(\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \in U_{n^*}] \\ &= \sum_{\sigma \in \mathcal{S}_{e^*}} f_{e^*}(\sigma)\mathbb{P}[C_\sigma(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \neq 1 |\, \sigma_{n^*} = \sigma] \\ &\quad \times \mathbb{P}[C_\sigma(\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \neq 1,\, \sigma^{-1}(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \neq \mathbf{y}^{\mathrm{IP}}_{\hat{\tau}^*-}(i),\, \sigma(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \neq \mathbf{y}^{\mathrm{IP}}_{\hat{\tau}^*-}(i) \,|\, \sigma_{n^*} = \sigma,\, C_\sigma(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \neq 1] \\ &\quad \times \frac{|U_{n^*}|}{(d-2)}\,\frac{(|U_{n^*}|-1)}{(d-3)} \\ &\geq \sum_{\sigma \in \mathcal{S}_{e^*}} f_{e^*}(\sigma)\left(\frac{d-N_\sigma(1)}{d}\right)\left(\frac{d-N_\sigma(1)-3}{d-1}\right)\left(\frac{d-5}{d-2}\right)\left(\frac{d-6}{d-3}\right)\end{aligned}$$

where in the last two lines we have used Part 1 of Assumption 1.2 (that f_e is constant on conjugacy classes). It follows (after some tedious arithmetic) that

$$\begin{aligned}&\mathbb{P}[\sigma_{n^*}^{-1}(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \in U_{n^*}\,,\, \sigma_{n^*}^{-1}(\mathbf{y}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \in U_{n^*}] \\ &\qquad \geq \frac{(d-5)(d-6)}{(d-1)(d-2)(d-3)}\big(d-3-(2d-5)\mathbb{P}[C_{\sigma_{n^*}}(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) = 1]\big) \\ &\qquad \geq \frac{(d-4)(d-5)(d-6)}{2(d-1)(d-2)(d-3)}\end{aligned} \tag{A.2}$$

where in the last step we have used the bound

$$\mathbb{P}[C_{\sigma_{n^*}}(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) = 1] \leq \mathbb{P}[C_{\sigma_{n^*}}(\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*-}(i)) \leq 3] \leq 1/4\,,$$

by Part 2 of Assumption 1.2. Since $d \geq 7$ the bound in (A.2) is at least $1/40$. Putting this into (A.1) we obtain

$$\mathbb{P}[\mathbf{x}^{\mathrm{IP}}_{\hat{\tau}^*}(i) = \mathbf{y}^{\mathrm{IP}}_{\hat{\tau}^*}(i)] \geq \frac{1}{80}.$$

We note that the above bound is also true conditional on the value of $\hat{\tau}^*$ and therefore since the coupling forces matched particles to move together once they have met, for any $\theta \in (0,1)$ such that $1/\theta$ is an integer,

$$\mathbb{P}[\mathbf{x}^{\mathrm{IP}}_{(i-1+\theta)\kappa T}(i) = \mathbf{y}^{\mathrm{IP}}_{(i-1+\theta)\kappa T}(i)] \geq \frac{1}{80}\mathbb{P}[\hat{\tau}^* < (i-1+\theta)\kappa T].$$

Furthermore, the construction implies that, viewed marginally, particles $\mathbf{x}(i)$ and $\mathbf{y}(i)$ behave as independent random walks until they are in the same edge which then rings in the Poisson process $\hat{\Lambda}$: this time is distributed as $M^{\mathrm{RW}}((\mathbf{x}(i), \mathbf{y}(i)))$. Thus

$$\mathbb{P}[\hat{\tau}^* \geq (i-1+\theta)\kappa T] \leq \sup_{u,v} \mathbb{P}[M^{\mathrm{RW}}(u,v) \geq \theta\kappa T] \leq 1/1000,$$

if $\theta\kappa \geq 10^{10}$, since G is easy. Thus

$$\mathbb{P}[\mathbf{x}^{\mathrm{IP}}_{(i-1+\theta)\kappa T}(i) = \mathbf{y}^{\mathrm{IP}}_{(i-1+\theta)\kappa T}(i)] \geq \frac{1}{80}(1 - 1/1000) > 1/100.$$

Now using the Markov property and the fact that once the two particles are matched they stay matched,

$$\mathbb{P}[\mathbf{x}^{\mathrm{IP}}_{i\kappa T}(i) = \mathbf{y}^{\mathrm{IP}}_{i\kappa T}(i)] \geq 1 - (99/100)^{1/\theta} > 1 - \delta$$

with $1/\theta = 100\lceil \log(\delta^{-1}) \rceil$. The proof is thus complete taking $\kappa = 10^{10}\theta^{-1} = 10^{12}\lceil \log(\delta^{-1}) \rceil$.

□

Proof of Lemma 3.7. Since the hypergraph is non-easy there exists $\mathbf{x} \in V^2$ such that

$$\mathbb{P}[M^{\mathrm{RW}}(\mathbf{x}) > 10^{10}T] > 1/1000.$$

We have

$$\begin{aligned}
&\mathbb{P}[M^{\mathrm{RW}}(\mathbf{x}) > 10^{10}T] \\
&= \mathbb{E}\big[\mathbb{E}[\mathbf{1}_{\{M^{\mathrm{RW}}(\mathbf{x}) > 10^{10}T\}} \big| \, \mathbf{x}^{\mathrm{RW}}_{(10^{10}-40)T}]\big] \\
&= \mathbb{E}\Big[\mathbb{E}\Big[\mathbf{1}_{\{M^{\mathrm{RW}}(\mathbf{x}) > (10^{10}-40)T\}} \big| \, \mathbf{x}^{\mathrm{RW}}_{(10^{10}-40)T}\Big] \\
&\quad \times \mathbb{E}\Big[\mathbf{1}_{\left\{M^{\mathrm{RW}}(\mathbf{x}^{\mathrm{RW}}_{(10^{10}-40)T}) > 40T\right\}} \big| \, \mathbf{x}^{\mathrm{RW}}_{(10^{10}-40)T}\Big]\Big] \\
&\leq \mathbb{E}\left[\mathbb{E}\Big[\mathbf{1}_{\{M^{\mathrm{RW}}(\mathbf{x}) > (10^{10}-40)T\}} \big| \, \mathbf{x}^{\mathrm{RW}}_{(10^{10}-40)T}\Big] \sup_{\mathbf{y} \in V^2} \mathbb{P}[M^{\mathrm{RW}}(\mathbf{y}) > 40T]\right] \\
&= \sup_{\mathbf{y} \in V^2} \mathbb{P}[M^{\mathrm{RW}}(\mathbf{y}) > 40T] \mathbb{P}\left[M^{\mathrm{RW}}(\mathbf{x}) > (10^{10}-40)T\right] \\
&\leq \left(\sup_{\mathbf{y} \in V^2} \mathbb{P}[M^{\mathrm{RW}}(\mathbf{y}) > 40T]\right)^{\frac{10^{10}}{40}}.
\end{aligned}$$

Hence there exists $\mathbf{y} \in V^2$ such that

$$\mathbb{P}[M^{\mathrm{RW}}(\mathbf{y}) > 40T] \geq \left(\frac{1}{1000}\right)^{40(10^{-10})} > 1 - 10^{-7}.$$

Now,

$$\begin{aligned}\mathbb{P}[M^{\mathrm{RW}}(\mathbf{y}) > 40T] &\leq \mathbb{P}\big[M^{\mathrm{RW}}(\mathbf{y}^{\mathrm{RW}}_{20T}) > 20T\big] \\ &= \sum_{\mathbf{v}\in V^2} \mathbb{P}[\mathbf{y}^{\mathrm{RW}}_{20T} = \mathbf{v}]\mathbb{P}\big[M^{\mathrm{RW}}(\mathbf{v}) > 20T\big].\end{aligned}$$

However, by Definition (2.2) of total-variation, Proposition 2.2, and the fact that (by the contraction principle) $T_{\mathrm{RW}(2,f,G)}(1/4) \leq T_{\mathrm{EX}(2,f,G)}(1/4)$,

$$\begin{aligned}\sum_{\mathbf{v}\in V^2} \mathbb{P}[\mathbf{y}^{\mathrm{RW}}_{20T} = \mathbf{v}]\mathbb{P}\big[M^{\mathrm{RW}}(\mathbf{v}) > 20T\big] - \sum_{\mathbf{v}\in V^2} \frac{\mathbb{P}\big[M^{\mathrm{RW}}(\mathbf{v}) > 20T\big]}{|V|^2} \\ \leq \|\mathcal{L}[\mathbf{y}^{\mathrm{RW}}_{20T}] - \mathrm{Unif}(V^2)\|_{\mathrm{TV}} \leq 2^{-20}.\end{aligned}$$

Hence

$$\begin{aligned}\sum_{\mathbf{v}\in V^2} \frac{\mathbb{P}\big[M^{\mathrm{RW}}(\mathbf{v}) > 20T\big]}{|V|^2} &\geq \sum_{\mathbf{v}\in V^2} \mathbb{P}[\mathbf{y}^{\mathrm{RW}}_{20T} = \mathbf{v}]\mathbb{P}\big[M^{\mathrm{RW}}(\mathbf{v}) > 20T\big] - 2^{-20} \\ &\geq 1 - 10^{-7} - 2^{-20} \\ &\geq 1 - \frac{1}{1000}.\end{aligned}$$

□

Proof of Lemma 3.8. We begin by conditioning on the value of $\mathbf{x}^{\mathrm{IP}}_{20T}$.

$$\mathbb{P}[\bar{M}^{\mathrm{RW}}(\mathbf{x}^{\mathrm{IP}}_{20T}) \leq 20T] = \sum_{(a,b,c,d)\in(V)_4} \mathbb{P}[\bar{M}^{\mathrm{RW}}((a,b,c,d)) \leq 20T]\mathbb{P}[\mathbf{x}^{\mathrm{IP}}_{20T} = (a,b,c,d)].$$

For each $\mathbf{a} \in (V)_4$, we have

$$\mathbb{P}[\bar{M}^{\mathrm{RW}}(\mathbf{a}) \leq 20T] \leq \sum_{i=1}^{4}\sum_{j=1}^{i-1} \mathbb{P}[M^{\mathrm{RW}}((\mathbf{a}(i),\mathbf{a}(j))) \leq 20T],$$

and so

$$\begin{aligned}&\mathbb{P}[\bar{M}^{\mathrm{RW}}(\mathbf{x}^{\mathrm{IP}}_{20T}) \leq 20T] \\ &\leq \sum_{i=1}^{4}\sum_{j=1}^{i-1}\sum_{\mathbf{u}\in(V)_2} \mathbb{P}[M^{\mathrm{RW}}(\mathbf{u}) \leq 20T]\mathbb{P}[(\mathbf{x}^{\mathrm{IP}}_{20T}(i),\mathbf{x}^{\mathrm{IP}}_{20T}(j)) = \mathbf{u}] \\ &\leq \sum_{i=1}^{4}\sum_{j=1}^{i-1}\sum_{\mathbf{u}\in(V)_2} \mathbb{P}[M^{\mathrm{RW}}(\mathbf{u}) \leq 20T]\mathbb{P}[\mathbf{x}^{\mathrm{IP}}_{20T}(i) \in \{\mathbf{u}(1),\mathbf{u}(2)\}, \mathbf{x}^{\mathrm{IP}}_{20T}(j) \in \{\mathbf{u}(1),\mathbf{u}(2)\}] \\ &= \sum_{i=1}^{4}\sum_{j=1}^{i-1}\sum_{\mathbf{u}\in(V)_2} \mathbb{P}[M^{\mathrm{RW}}(\mathbf{u}) \leq 20T]\mathbb{P}[\{\mathbf{x}(i),\mathbf{x}(j)\}^{\mathrm{EX}}_{20T} = \{\mathbf{u}(1),\mathbf{u}(2)\}] \\ &= 2\sum_{i=1}^{4}\sum_{j=1}^{i-1}\sum_{\mathbf{u}\in\binom{V}{2}} \mathbb{P}[M^{\mathrm{RW}}(\mathbf{u}) \leq 20T]\mathbb{P}[\{\mathbf{x}(i),\mathbf{x}(j)\}^{\mathrm{EX}}_{20T} = \{\mathbf{u}(1),\mathbf{u}(2)\}]. \qquad \text{(A.3)}\end{aligned}$$

Now, for each $1 \le j < i \le 4$ (and motivated by Oliveira (2013)) set

$$\mathrm{Good}_{i,j} := \left\{ \{a,b\} \in \binom{V}{2} : \left| \mathbb{P}[\{\mathbf{x}(i),\mathbf{x}(j)\}^{\mathrm{EX}}_{20T} = \{a,b\}] - \frac{1}{\binom{|V|}{2}} \right| \le \frac{\varepsilon}{\binom{|V|}{2}} \right\}.$$

We decompose the sum over $\mathbf{u}$ above into $\mathbf{u} \in \mathrm{Good}_{i,j}$ and $\mathbf{u} \in \mathrm{Bad}_{i,j}$, where

$$\mathrm{Bad}_{i,j} = \binom{V}{2} \setminus \mathrm{Good}_{i,j}.$$

For the Good terms, we have

$$\begin{aligned}
&2\sum_{i=1}^{4}\sum_{j=1}^{i-1}\sum_{\mathbf{u}\in\mathrm{Good}_{i,j}} \mathbb{P}[M^{\mathrm{RW}}(\mathbf{u}) \le 20T]\mathbb{P}[\{\mathbf{x}(i),\mathbf{x}(j)\}^{\mathrm{EX}}_{20T} = \{\mathbf{u}(1),\mathbf{u}(2)\}] \\
&\le 2\sum_{i=1}^{4}\sum_{j=1}^{i-1}\sum_{\mathbf{u}\in\mathrm{Good}_{i,j}} \mathbb{P}[M^{\mathrm{RW}}(\mathbf{u}) \le 20T]\frac{(1+\varepsilon)}{\binom{|V|}{2}} \\
&\le 25(1+\varepsilon)\sum_{\mathbf{u}\in V^2}\frac{\mathbb{P}[M^{\mathrm{RW}}(\mathbf{u}) \le 20T]}{|V|^2}. \qquad \text{(A.4)}
\end{aligned}$$

For the Bad terms, we have

$$\begin{aligned}
&2\sum_{i=1}^{4}\sum_{j=1}^{i-1}\sum_{\mathbf{u}\in\mathrm{Bad}_{i,j}} \mathbb{P}[M^{\mathrm{RW}}(\mathbf{u}) \le 20T]\mathbb{P}[\{\mathbf{x}(i),\mathbf{x}(j)\}^{\mathrm{EX}}_{20T} = \{\mathbf{u}(1),\mathbf{u}(2)\}] \\
&\le 2\sum_{i=1}^{4}\sum_{j=1}^{i-1}\sum_{\mathbf{u}\in\mathrm{Bad}_{i,j}} \mathbb{P}[\{\mathbf{x}(i),\mathbf{x}(j)\}^{\mathrm{EX}}_{20T} = \{\mathbf{u}(1),\mathbf{u}(2)\}] \\
&= 2\sum_{i=1}^{4}\sum_{j=1}^{i-1}\mathbb{P}[\{\mathbf{x}(i),\mathbf{x}(j)\}^{\mathrm{EX}}_{20T} \in \mathrm{Bad}_{i,j}]. \qquad \text{(A.5)}
\end{aligned}$$

Note that for each $1 \le j < i \le 4$,

$$\frac{1}{2}\sum_{\mathbf{u}\in\binom{V}{2}} \left| \mathbb{P}[\{\mathbf{x}(i),\mathbf{x}(j)\}^{\mathrm{EX}}_{20T} = \{\mathbf{u}(1),\mathbf{u}(2)\}] - \frac{1}{\binom{|V|}{2}} \right| > \frac{1}{2}\frac{\varepsilon}{\binom{|V|}{2}}|\mathrm{Bad}_{i,j}|,$$

since every $\mathbf{u} \in \mathrm{Bad}_{i,j}$ contributes at least $\varepsilon/\binom{|V|}{2}$ to the sum. However, the left-hand side in the above equation is at most 2^{-20} by the choice of T and Definition (2.3) of total-variation. We deduce that

$$|\mathrm{Bad}_{i,j}| \le \varepsilon^{-1}2^{-19}\binom{|V|}{2},$$

and thus

$$|\mathrm{Good}_{i,j}| \ge (1-\varepsilon^{-1}2^{-19})\binom{|V|}{2}.$$

However, for each $\mathbf{u} \in \mathrm{Good}_{i,j}$ we know that

$$\mathbb{P}[\{\mathbf{x}(i),\mathbf{x}(j)\}^{\mathrm{EX}}_{20T} = \{\mathbf{u}(1),\mathbf{u}(2)\}] \ge \frac{1-\varepsilon}{\binom{|V|}{2}}.$$

Therefore,

$$\mathbb{P}[\{\mathbf{x}(i),\mathbf{x}(j)\}^{\mathrm{EX}}_{20T} \in \mathrm{Good}_{i,j}\}] \geq \frac{1-\varepsilon}{\binom{|V|}{2}}|\mathrm{Good}_{i,j}| \geq 1-\varepsilon-\varepsilon^{-1}2^{-19}.$$

We deduce that

$$\mathbb{P}[\{\mathbf{x}(i),\mathbf{x}(j)\}^{\mathrm{EX}}_{20T} \in \mathrm{Bad}_{i,j}\}] \leq \varepsilon+\varepsilon^{-1}2^{-19}.$$

Plugging this into (A.5) gives

$$\begin{aligned}2\sum_{i=1}^{4}\sum_{j=1}^{i-1}\sum_{\mathbf{u}\in\mathrm{Bad}_{i,j}} &\mathbb{P}[M^{\mathrm{RW}}(\mathbf{u}) \leq 20T]\mathbb{P}[\{\mathbf{x}(i),\mathbf{x}(j)\}^{\mathrm{EX}}_{20T} = \{\mathbf{u}(1),\mathbf{u}(2)\}] \\ &\leq 12(\varepsilon+\varepsilon^{-1}2^{-19}).\end{aligned}$$

Combining this with (A.4) and (A.3) gives

$$\mathbb{P}[\bar{M}^{\mathrm{RW}}(\mathbf{x}^{\mathrm{IP}}_{20T}) \leq 20T] \leq 12(\varepsilon+\varepsilon^{-1}2^{-19}) + 25(1+\varepsilon)\sum_{\mathbf{u}\in V^2}\frac{\mathbb{P}[M^{\mathrm{RW}}(\mathbf{u}) \leq 20T]}{|V|^2}. \qquad \square$$

Proof of Proposition 3.9. This proof is similar to the proof of Proposition 4.6 in Oliveira (2013). We present a coupling of $\{\mathbf{x}^{\mathrm{IP}}_t\}_{t\geq 0}$ and $\{\mathbf{x}^{\mathrm{RW}}_t\}_{t\geq 0}$ such that the two processes agree up to time $\bar{M}^{\mathrm{RW}}(\mathbf{x})$. The coupling has state-space $S := (V)_2 \times V^2$ which we split into two parts: $\Delta := \{(\mathbf{z},\mathbf{z}) : \mathbf{z} \in (V)_2\}$ and $\Delta^{\complement}$. Denote by $q(\cdot,\cdot)$ the transition rates. We construct the coupling as follows:

1. if $(\mathbf{x},\mathbf{y}) \in \Delta^{\complement}$, the transition rate to any other state in S is the same as that of independent realisations of $\mathrm{IP}(4,f,G)$ and $\mathrm{RW}(4,f,G)$.
2. if $(\mathbf{x},\mathbf{x}) \in \Delta$,
   (a) for $e \in E$ with $|e \cap \{x(1),x(2),x(3),x(4)\}| = 1$ and for each $\sigma_e \in \mathcal{S}_e$,
   $$q\big((\mathbf{x},\mathbf{x}),(\sigma_e(\mathbf{x}),\sigma_e(\mathbf{x}))\big) = f_e(\sigma_e).$$
   (b) for $e \in E$ with $|e \cap \{\mathbf{x}(1),\mathbf{x}(2),\mathbf{x}(3),\mathbf{x}(4)\}| > 1$ and for each $\sigma_e \in \mathcal{S}_e$,
   $$\begin{aligned}
   q\Big(\big(\mathbf{x},\mathbf{x}\big),\big(\sigma_e(\mathbf{x}),(\sigma_e(\mathbf{x}(1)),\mathbf{x}(2),\mathbf{x}(3),\mathbf{x}(4))\big)\Big) &= f_e(\sigma_e),\\
   q\Big(\big(\mathbf{x},\mathbf{x}\big),\big(\mathbf{x},(\mathbf{x}(1),\sigma_e(\mathbf{x}(2)),\mathbf{x}(3),\mathbf{x}(4))\big)\Big) &= f_e(\sigma_e),\\
   q\Big(\big(\mathbf{x},\mathbf{x}\big),\big(\mathbf{x},(\mathbf{x}(1),\mathbf{x}(2),\sigma_e(\mathbf{x}(3)),\mathbf{x}(4))\big)\Big) &= f_e(\sigma_e),\\
   q\Big(\big(\mathbf{x},\mathbf{x}\big),\big(\mathbf{x},(\mathbf{x}(1),\mathbf{x}(2),\mathbf{x}(3),\sigma_e(\mathbf{x}(4)))\big)\Big) &= f_e(\sigma_e).
   \end{aligned}$$
3. all other transitions have rate 0.

By inspection of the marginals it is clear that this indeed gives a coupling of the two processes. Furthermore, if we start the coupling from a state $\mathbf{x} \in \Delta$, the two processes can only differ after a transition has occurred according to rule 2.(b); but the first time this happens is precisely $\bar{M}^{\mathrm{RW}}(\mathbf{x})$. $\square$

B Technical proofs for Section 6

Here we include proofs of some of the results used when counting the number of pinkenings in the chameleon process.

Proof of Lemma 6.1. Fix $b \in I(W^{\mathfrak{C}}) = (W^{\mathfrak{C}})_T$ (as $[0,T]$ is a constant-colour phase) and suppose that there exists $a \in I(R)$ with $\phi_a = \phi_b$ and $F_a = (b,c,d)$ for some $c,d \in I(W^r)$. It is trivial that $I_{[T,\phi_a)}(b) \in W_{\phi_a-}$ since the particle at b at time T cannot be pinkened if it has not been in any ringing edge. Similarly we have $I_{[T,\phi_a)}(a) \in R_{\phi_a-}$. Since particles at c and d at time T cannot be pinkened under the modified chameleon process, we must have $I_{[T,\phi_a)}(c), I_{[T,\phi_a)}(d) \in W_{\phi_a-}$. Since $F_a = (b,c,d)$ we know that there exists $j \in \{1,\dots,d'(c_{t_a})\}$ and we are in one of 4 situations, corresponding to the 4 cases listed just before Lemma 6.1:

(i). $a = c_{t_a}^{2j-2}(m(c_{t_a}))$ and

$$I_{[T,\phi_a)}(b,c,d) = \big(c_{t_a}^{2d'(c_{t_a})+2j-2}(m(c_{t_a})), c_{t_a}^{2j-1}(m(c_{t_a})), c_{t_a}^{2d'(c_{t_a})+2j-1}(m(c_{t_a}))\big),$$

with $j \in B_i^2(R_{\phi_a-}, W_{\phi_a-}, (\sigma_{\phi_a})_i) \cap \tilde{D}_i^3(R_{\phi_a-}, W_{\phi_a-}, (\sigma_{\phi_a})_i)$ for some $i \in \{1,\dots,K(\sigma_{\phi_a})\}$ which implies $j \in \tilde{C}_i^1(R_{\phi_a-}, W_{\phi_a-}, (\sigma_{\phi_a})_i)$.

(ii). $a = c_{t_a}^{2d'(c_{t_a})+2j-2}(m(c_{t_a}))$ and

$$I_{[T,\phi_a)}(b,c,d) = \big(c_{t_a}^{2j-2}(m(c_{t_a})), c_{t_a}^{2j-1}(m(c_{t_a})), c_{t_a}^{2d'(c_{t_a})+2j-1}(m(c_{t_a}))\big),$$

with $j \in B_i^1(R_{\phi_a-}, W_{\phi_a-}, (\sigma_{\phi_a})_i) \cap \tilde{D}_i^3(R_{\phi_a-}, W_{\phi_a-}, (\sigma_{\phi_a})_i)$ for some $i \in \{1,\dots,K(\sigma_{\phi_a})\}$ which implies $j \in \tilde{C}_i^1(R_{\phi_a-}, W_{\phi_a-}, (\sigma_{\phi_a})_i)$.

(iii). $a = c_{t_a}^{2j-1}(m(c_{t_a}))$ and

$$I_{[T,\phi_a)}(b,c,d) = \big(c_{t_a}^{2d'(c_{t_a})+2j-1}(m(c_{t_a})), c_{t_a}^{2j-2}(m(c_{t_a})), c_{t_a}^{2d'(c_{t_a})+2j-2}(m(c_{t_a}))\big),$$

with $j \in \tilde{B}_i^4(R_{\phi_a-}, W_{\phi_a-}, (\sigma_{\phi_a})_i) \cap \tilde{D}_i^2(R_{\phi_a-}, W_{\phi_a-}, (\sigma_{\phi_a})_i)$ for some $i \in \{1,\dots,K(\sigma_{\phi_a})\}$ which implies $j \in \tilde{C}_i^2(R_{\phi_a-}, W_{\phi_a-}, (\sigma_{\phi_a})_i)$.

(iv). $a = c_{t_a}^{2d'(c_{t_a})+2j-1}(m(c_{t_a}))$ and

$$I_{[T,\phi_a)}(b,c,d) = \big(c_{t_a}^{2j-1}(m(c_{t_a})), c_{t_a}^{2j-2}(m(c_{t_a})), c_{t_a}^{2d'(c_{t_a})+2j-2}(m(c_{t_a}))\big),$$

with $j \in \tilde{B}_i^4(R_{\phi_a-}, W_{\phi_a-}, (\sigma_{\phi_a})_i) \cap \tilde{D}_i^1(R_{\phi_a-}, W_{\phi_a-}, (\sigma_{\phi_a})_i)$ for some $i \in \{1,\dots,K(\sigma_{\phi_a})\}$ which implies $j \in \tilde{C}_i^2(R_{\phi_a-}, W_{\phi_a-}, (\sigma_{\phi_a})_i)$.

In all four cases there will be a pinkening at vertices a and $I_{[T,\phi_a)}(b)$ at time ϕ_a. □

Remark B.1. The proof of Lemma 6.1 reveals the reason behind reserving some of the white particles. Indeed, if the particles at c and d at time T were able to change colour then it may be that by time ϕ_a- they are no longer white. In this case, we would not be able to have a pinkening at time ϕ_a.

Proof of Lemma 6.2. For a fixed $a \in V$ we wish to upper bound $\mathbb{P}[F_a = *|\, \phi_a \neq \infty]$. On the event $\{\phi_a \neq *\}$, the event $\{F_a = *\}$ is equivalent to the event that at time ϕ_a one of the following occurs:

Event H_1: $|c_{t_a}| < 4$,
Event H_2: $|c_{t_a}| \geq 4$ but there does not exist a $j \in \{1, \dots, d'(c_{t_a})\}$ and
$k \in \{2j-2, 2d'(c_{t_a}) + 2j - 2, 2j-1, 2d'(c_{t_a}) + 2j - 1\}$ with $a = c_{t_a}^k(m(c_{t_a}))$.

We have

$$\mathbb{P}[F_a = *|\, \phi_a \neq \infty] = \mathbb{P}[H_1] + \mathbb{P}[H_2|H_1']\mathbb{P}[H_1'] \,.$$

Now,

$$\mathbb{P}[H_2|c_{t_a}] = \frac{|c_{t_a}| \bmod 4}{|c_{t_a}|}$$

and so $\mathbb{P}[H_2|H_1'] \leq 3/7$. Furthermore, $\mathbb{P}[H_1] \leq 1/4$ by Assumption 1.2. Thus $\mathbb{P}[F_a = *|\, \phi_a \neq \infty] \leq 4/7$, as required. □

Proof of Proposition 6.3. We imitate the proof of Proposition 9.1 of Oliveira (2013). Let $\mathbf{x}_t^{\mathrm{IP}}$ be a realisation of $\mathrm{IP}(4, f, G)$ with $\mathbf{x}_0^{\mathrm{IP}} = (a, b, c, d)$. Then

$$\begin{aligned}
\mathbb{P}[a \in I(R), b \in I(W^{\complement}), c, d \in I(W^r)] &= \mathbb{P}[\mathbf{x}_T^{\mathrm{IP}} \in R \times W^{\complement} \times W^r \times W^r] \\
&\geq (1 - 2^{-9})^2 \frac{|R \times W^{\complement} \times W^r \times W^r|}{|(V)_4|} \\
&\geq (1 - 2^{-9})^2 \frac{|R||W^{\complement}||W^r|^2}{|V|^4} \,,
\end{aligned}$$

where the first inequality follows from Propositions 2.2 and 2.4, thanks to the supposed lower bound on T. □

References

Aldous, D. and J. A. Fill (2002). Reversible markov chains and random walks on graphs. Unfinished monograph, recompiled 2014, available at `http://www.stat.berkeley.edu/~aldous/RWG/book.html`.

Berestycki, N., O. Schramm, and O. Zeitouni (2011). Mixing times for random k-cycles and coalescence-fragmentation chains. *Ann. Probab. 39*(5), 1815–1843.

Berestycki, N. and B. Sengul (2014). Cutoff for conjugacy-invariant random walks on the permutation group. *Preprint available at http://arxiv.org/abs/1410.4800*.

Caputo, P., T. M. Liggett, and T. Richthammer (2010). Proof of Aldous' spectral gap conjecture. *J. Amer. Math. Soc. 23*(3), 831–851.

Chung, F. and A. Tsiatas (2014). Hypergraph coloring games and voter models. *Internet Math. 10*(1-2), 66–86.

Diaconis, P. (1988). *Group representations in probability and statistics*. Institute of Mathematical Statistics Lecture Notes—Monograph Series, 11. Institute of Mathematical Statistics, Hayward, CA.

Diaconis, P. and L. Saloff-Coste (1993). Comparison techniques for random walk on finite groups. *Ann. Probab. 21*(4), 2131–2156.

Diaconis, P. and M. Shahshahani (1981). Generating a random permutation with random transpositions. *Z. Wahrsch. Verw. Gebiete 57*(2), 159–179.

Erikshed, S. (2011). Mixing times for neighbour transposition shuffles on graphs. *Preprint available at http://www.dissertations.se/dissertation/7ca577aafc/*.

Flatto, L., A. M. Odlyzko, and D. B. Wales (1985). Random shuffles and group representations. *Ann. Probab. 13*(1), 154–178.

Hough, B. (2015). The random k cycle walk on the symmetric group. *Probability Theory and Related Fields*, 1–36.

Istrate, G., C. Bonchis, and M. Marin (2014). Interactive particle systems and random walks on hypergraphs.

Jonasson, J. (2012). Mixing times for the interchange process. *ALEA Lat. Am. J. Probab. Math. Stat. 9*(2), 667–683.

Lacoin, H. (2016). Mixing time and cutoff for the adjacent transposition shuffle and the simple exclusion. *Ann. Probab. 44*(2), 1426–1487.

Levin, D. A., Y. Peres, and E. L. Wilmer (2008). *Markov Chains and Mixing Times*. American Mathematical Society.

Liggett, T. M. (1985). *Interacting particle systems*, Volume 276 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, New York.

Liggett, T. M. (1999). *Stochastic interacting systems: contact, voter and exclusion processes*, Volume 324 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin.

Lulov, N. and I. Pak (2002). Rapidly mixing random walks and bounds on characters of the symmetric group. *J. Algebraic Combin. 16*(2), 151–163.

Morris, B. (2006). The mixing time for simple exclusion. *The Annals of Applied Probability 16*(2), 615–635.

Oliveira, R. I. (2013). Mixing of the symmetric exclusion processes in terms of the corresponding single-particle random walk. *The Annals of Probability 41*(2), 871–913.

Roichman, Y. (1996). Upper bound on the characters of the symmetric groups. *Invent. Math. 125*(3), 451–485.

SBC: Department of Mathematics, University of York, York, YO10 5DD, UK.
Email: `stephen.connor@york.ac.uk`
WWW: `maths.york.ac.uk/www/sbc502`

RJP: Department of Statistics, University College London, London, WC1E 6BT, UK.
Email: `r.pymar@ucl.ac.uk`
WWW: `www.homepages.ucl.ac.uk/~ucahrjp`